CISSP®
PASSPORT

FIRST EDITION

- Complete coverage of all current exam objectives for the 2021 Certified Information Systems Security Professional exam
- 300 practice exam questions

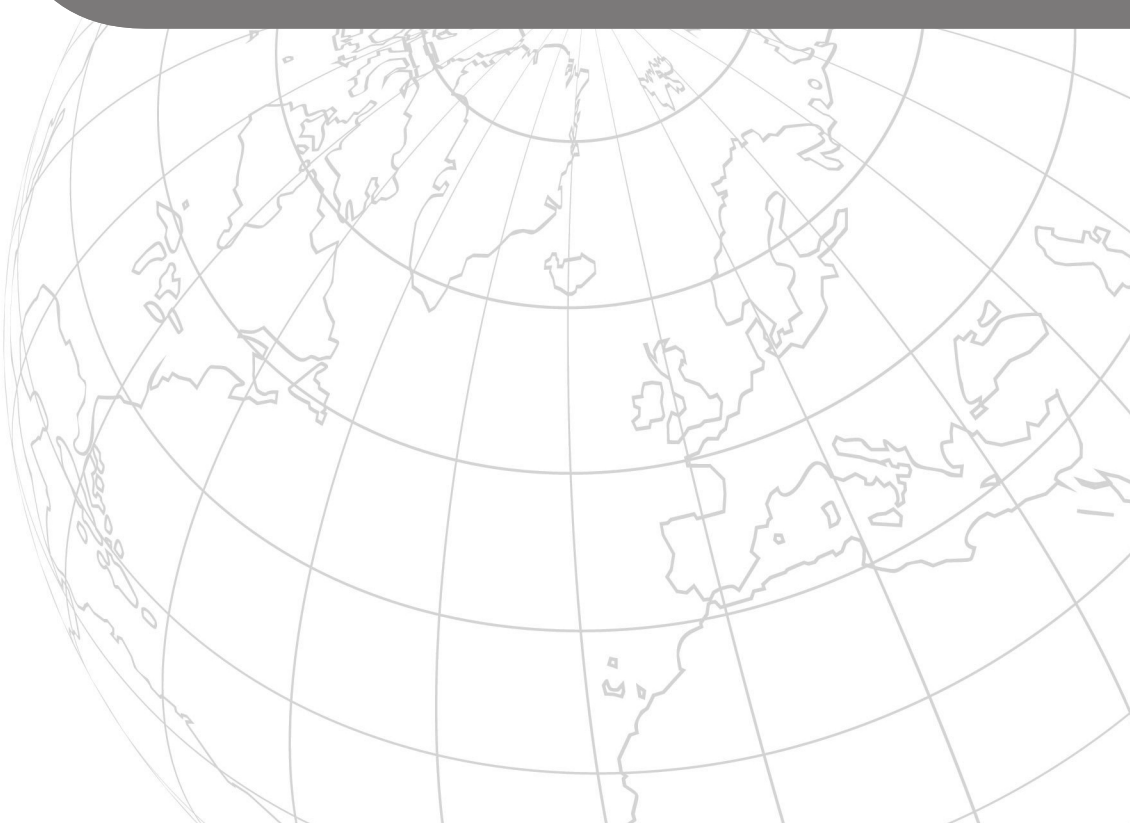
BOBBY E. ROGERS

CISSP-ISSEP®, CRISC®, CEH™, MCSE: Security, CompTIA A+™, Network+™, and Security+™

Mc Graw Hill

ACCELERATED REVIEW FOR QUICK EXAM PREPARATION

CISSP®
PASSPORT



About the Author

Bobby Rogers (he/his/him) is a cybersecurity professional with over 30 years in the information technology and cybersecurity fields. He currently works with a major engineering company in Huntsville, Alabama, helping to secure networks and manage cyber risk for its customers. Bobby's customers include the U.S. Army, NASA, the State of Tennessee, and private/commercial companies and organizations. His specialties are cybersecurity engineering, security compliance, and cyber risk management, but he has worked in almost every area of cybersecurity, including network defense, computer forensics and incident response, and penetration testing.

Bobby is a retired Master Sergeant from the U.S. Air Force, having served for over 21 years. He has built and secured networks in the United States, Chad, Uganda, South Africa, Germany, Saudi Arabia, Pakistan, Afghanistan, and several other remote locations. His decorations include two Meritorious Service medals, three Air Force Commendation medals, the National Defense Service medal, and several Air Force Achievement medals. He retired from active duty in 2006.

Bobby has a master of science in information assurance and a bachelor of science in computer information systems (with a dual concentration in Russian language), and two associate of science degrees. His many certifications include CISSP-ISSEP, CRISC, CySA+, CEH, and MCSE: Security.

Bobby has narrated and produced over 30 computer training videos for several training companies and currently produces them for Pluralsight (https://www.pluralsight.com). He is also the author of *CompTIA Mobility+ All-in-One Exam Guide (Exam MB0-001)*, *CRISC Certified in Risk and Information Systems Control All-in-One Exam Guide*, and *Mike Meyers' CompTIA Security+ Certification Guide (Exam SY0-401)*, and is the contributing author/technical editor for the popular *CISSP All-in-One Exam Guide, Ninth Edition*, all of which are published by McGraw Hill.

About the Technical Editor

Nichole O'Brien is a creative business leader with over 25 years of experience in cybersecurity and IT leadership, program management, and business development across commercial, education, and federal, state, and local business markets. Her focus on innovative solutions is demonstrated by the development of a commercial cybersecurity and IT business group, which she currently manages in a Fortune 500 corporation and has received the corporation's annual Outstanding Customer Service Award. She currently serves as Vice President of Outreach for Cyber Huntsville, is on the Foundation Board for the National Cyber Summit, and supports cyber education initiatives like the USSRC Cyber Camp. Nichole has bachelor's and master's degrees in business administration and has a CISSP certification.

CISSP®
PASSPORT

Bobby E. Rogers



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

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

Copyright © 2023 by McGraw Hill. All rights reserved. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

ISBN: 978-1-26-427798-8
MHID: 1-26-427798-9

The material in this eBook also appears in the print version of this title: ISBN: 978-1-26-427797-1,
MHID: 1-26-427797-0.

eBook conversion by codeMantra
Version 1.0

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

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

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

TERMS OF USE

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

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

I'd like to dedicate this book to the cybersecurity professionals who tirelessly, and sometimes, thanklessly, protect our information and systems from all who would do them harm.
I also dedicate this book to the people who serve in uniform as military personnel, public safety professionals, police, firefighters, and medical professionals, sacrificing sometimes all that they are and have so that we may all live in peace, security, and safety.

—Bobby Rogers

This page intentionally left blank

Contents at a Glance

1.0 Security and Risk Management 1
2.0 Asset Security 85
3.0 Security Architecture and Engineering 115
4.0 Communication and Network Security 183
5.0 Identity and Access Management (IAM) 225
6.0 Security Assessment and Testing 259
7.0 Security Operations 285
8.0 Software Development Security 393
A About the Online Content 427

Index 431

This page intentionally left blank

Contents

Acknowledgments xxvii
Introduction xxix

1.0 Security and Risk Management 1

Objective 1.1 Understand, adhere to, and promote professional ethics 2
The (ISC)² Code of Ethics 3
Code of Ethics Preamble 3
Code of Ethics Canons 3
Organizational Code of Ethics 4
Workplace Ethics Statements and Policies 4
Other Sources for Ethics Requirements 5
REVIEW 7
1.1 QUESTIONS 7
1.1 ANSWERS 8

Objective 1.2 Understand and apply security concepts 9
Security Concepts 9
Data, Information, Systems, and Entities 9
Confidentiality 10
Integrity 11
Availability 11
Supporting Tenets of Information Security 11
Identification 11
Authentication 11
Authenticity 12
Authorization 12
Auditing and Accountability 12
Nonrepudiation 12
Supporting Security Concepts 13

REVIEW 14
1.2 QUESTIONS 14
1.2 ANSWERS 15
Objective 1.3 Evaluate and apply security governance principles . . . 16
Security Governance 16
External Governance 16
Internal Governance 16
Alignment of Security Functions to Business Requirements 17
Business Strategy and Security Strategy 17
Organizational Processes 18
Organizational Roles and Responsibilities 18
Security Control Frameworks 19
Due Care/Due Diligence 20
REVIEW 21
1.3 QUESTIONS 21
1.3 ANSWERS 22
Objective 1.4 Determine compliance and other requirements 23
Compliance 23
Legal and Regulatory Compliance 24
Contractual Compliance 25
Compliance with Industry Standards 25
Privacy Requirements 25
REVIEW 26
1.4 QUESTIONS 27
1.4 ANSWERS 28
Objective 1.5 Understand legal and regulatory issues that pertain to information security in a holistic context 29
Legal and Regulatory Requirements 29
Cybercrimes 29
Licensing and Intellectual Property Requirements 30
Import/Export Controls 31
Transborder Data Flow 32
Privacy Issues 32
REVIEW 33
1.5 QUESTIONS 33
1.5 ANSWERS 34
Objective 1.6 Understand requirements for investigation types (i.e., administrative, criminal, civil, regulatory, industry standards) . . . 35
Investigations 35
Administrative Investigations 35
Civil Investigations 35

Criminal Investigations 36
Regulatory Investigations 36
Industry Standards for Investigations 37
REVIEW 37
1.6 QUESTIONS 38
1.6 ANSWERS 39
Objective 1.7 Develop, document, and implement security policy, standards, procedures, and guidelines 39
Internal Governance 40
Policy 40
Procedures 40
Standards 41
Guidelines 41
Baselines 42
REVIEW 42
1.7 QUESTIONS 43
1.7 ANSWERS 44
Objective 1.8 Identify, analyze, and prioritize Business Continuity (BC) requirements 45
Business Continuity 45
Business Impact Analysis 46
Developing the BIA 46
REVIEW 47
1.8 QUESTIONS 47
1.8 ANSWERS 48
Objective 1.9 Contribute to and enforce personnel security policies and procedures 48
Personnel Security 49
Candidate Screening and Hiring 49
Employment Agreements and Policies 50
Onboarding, Transfers, and Termination Processes 50
Vendor, Consultant, and Contractor Agreements and Controls 52
Compliance Policy Requirements 53
Privacy Policy Requirements 53
REVIEW 54
1.9 QUESTIONS 55
1.9 ANSWERS 56
Objective 1.10 Understand and apply risk management concepts . . . 57
Risk Management 57
Elements of Risk 57
Identify Threats and Vulnerabilities 59

Risk Assessment/Analysis 60
Risk Response 63
Risk Frameworks 64
Countermeasure Selection and Implementation 64
Applicable Types of Controls 65
Control Assessments (Security and Privacy) 66
Monitoring and Measurement 67
Reporting 67
Continuous Improvement 68
REVIEW 68
1.10 QUESTIONS 69
1.10 ANSWERS 69
Objective 1.11 Understand and apply threat modeling concepts and methodologies 70
Threat Modeling 70
Threat Components 70
Threat Modeling Methodologies 72
REVIEW 73
1.11 QUESTIONS 73
1.11 ANSWERS 73
Objective 1.12 Apply Supply Chain Risk Management (SCRM) concepts 74
Supply Chain Risk Management 74
Risks Associated with Hardware, Software, and Services 74
Third-Party Assessment and Monitoring 76
Minimum Security Requirements 77
Service Level Requirements 77
REVIEW 77
1.12 QUESTIONS 78
1.12 ANSWERS 79
Objective 1.13 Establish and maintain a security awareness, education, and training program. 80
Security Awareness, Education, and Training Program 80
Methods and Techniques to Present Awareness and Training 80
Periodic Content Reviews 82
Program Effectiveness Evaluation 82
REVIEW 82
1.13 QUESTIONS 83
1.13 ANSWERS 84

2.0 Asset Security 85

Objective 2.1 Identify and classify information and assets. 86
Asset Classification 86
Data Classification 87
REVIEW 89
2.1 QUESTIONS 89
2.1 ANSWERS 90

Objective 2.2 Establish information and asset handling requirements 90
Information and Asset Handling 90
Handling Requirements 91
Information Classification and Handling Systems 93
REVIEW 94
2.2 QUESTIONS 95
2.2 ANSWERS 95

Objective 2.3 Provision resources securely 96
Securing Resources 96
Asset Ownership 96
Asset Inventory 96
Asset Management 97
REVIEW 98
2.3 QUESTIONS 99
2.3 ANSWERS 99

Objective 2.4 Manage data lifecycle 99
Managing the Data Life Cycle 100
Data Roles 100
Data Collection 102
Data Location 102
Data Maintenance 102
Data Retention 103
Data Remanence 103
Data Destruction 103
REVIEW 104
2.4 QUESTIONS 104
2.4 ANSWERS 105

Objective 2.5 Ensure appropriate asset retention (e.g., End-of-Life (EOL), End-of-Support (EOS)). 105
Asset Retention 105
Asset Life Cycle 106
End-of-Life and End-of-Support 106

REVIEW . . . 108
2.5 QUESTIONS . . . 108
2.5 ANSWERS . . . 108
Objective 2.6 Determine data security controls and compliance requirements . . . 109
Data Security and Compliance . . . 109
Data States . . . 109
Control Standards Selection . . . 110
Scoping and Tailoring Data Security Controls . . . 111
Data Protection Methods . . . 111
REVIEW . . . 113
2.6 QUESTIONS . . . 113
2.6 ANSWERS . . . 114
3.0 Security Architecture and Engineering . . . 115
Objective 3.1 Research, implement, and manage engineering processes using secure design principles . . . 116
Threat Modeling . . . 116
Least Privilege . . . 116
Defense in Depth . . . 117
Secure Defaults . . . 117
Fail Securely . . . 117
Separation of Duties . . . 118
Keep It Simple . . . 119
Zero Trust . . . 119
Privacy by Design . . . 119
Trust But Verify . . . 119
Shared Responsibility . . . 120
REVIEW . . . 120
3.1 QUESTIONS . . . 121
3.1 ANSWERS . . . 122
Objective 3.2 Understand the fundamental concepts of security models (e.g., Biba, Star Model, Bell-LaPadula) . . . 122
Security Models . . . 122
Terms and Concepts . . . 123
System States and Processing Modes . . . 124
Confidentiality Models . . . 126
Integrity Models . . . 127
Other Access Control Models . . . 128
REVIEW . . . 128
3.2 QUESTIONS . . . 129
3.2 ANSWERS . . . 130

Objective 3.3 Select controls based upon systems security requirements 130
Selecting Security Controls 130
Performance and Functional Requirements 131
Data Protection Requirements 131
Governance Requirements 132
Interface Requirements 132
Risk Response Requirements 133
REVIEW 133
3.3 QUESTIONS 134
3.3 ANSWERS 134
Objective 3.4 Understand security capabilities of Information Systems (IS) (e.g., memory protection, Trusted Platform Module (TPM), encryption/decryption) 135
Information System Security Capabilities 135
Hardware and Firmware System Security 135
Secure Processing 137
REVIEW 138
3.4 QUESTIONS 139
3.4 ANSWERS 139
Objective 3.5 Assess and mitigate the vulnerabilities of security architectures, designs, and solution elements 139
Vulnerabilities of Security Architectures, Designs, and Solutions 140
Client-Based Systems 140
Server-Based Systems 140
Distributed Systems 141
Database Systems 141
Cryptographic Systems 142
Industrial Control Systems 142
Internet of Things 143
Embedded Systems 143
Cloud-Based Systems 144
Virtualized Systems 145
Containerization 146
Microservices 146
Serverless 146
High-Performance Computing Systems 146
Edge Computing Systems 146
REVIEW 147
3.5 QUESTIONS 148
3.5 ANSWERS 148

Objective 3.6 Select and determine cryptographic solutions 148
Cryptography 149
Cryptographic Life Cycle 149
Cryptographic Methods 151
Integrity 154
Hybrid Cryptography 155
Digital Certificates 156
Public Key Infrastructure 156
Nonrepudiation and Digital Signatures 158
Key Management Practices 158
REVIEW 159
3.6 QUESTIONS 160
3.6 ANSWERS 161
Objective 3.7 Understand methods of cryptanalytic attacks. 161
Cryptanalytic Attacks 161
Brute Force 162
Ciphertext Only 162
Known Plaintext 162
Chosen Ciphertext and Chosen Plaintext 163
Frequency Analysis 163
Implementation 163
Side Channel 163
Fault Injection 164
Timing 164
Man-in-the-Middle (On-Path) 164
Pass the Hash 165
Kerberos Exploitation 165
Ransomware 165
REVIEW 166
3.7 QUESTIONS 166
3.7 ANSWERS 167
Objective 3.8 Apply security principles to site and facility design . . . 167
Site and Facility Design 167
Site Planning 167
Secure Design Principles 168
REVIEW 172
3.8 QUESTIONS 172
3.8 ANSWERS 173
Objective 3.9 Design site and facility security controls. 173
Designing Facility Security Controls 173
Crime Prevention Through Environmental Design 174
Key Facility Areas of Concern 174

REVIEW . . . 181
3.9 QUESTIONS . . . 181
3.9 ANSWERS . . . 182

4.0 Communication and Network Security . . . 183

Objective 4.1 Assess and implement secure design principles in network architectures . . . 184
Fundamental Networking Concepts . . . 184
Open Systems Interconnection and Transmission Control Protocol/Internet Protocol Models . . . 185
Internet Protocol Networking . . . 187
Secure Protocols . . . 189
Application of Secure Networking Concepts . . . 193
Implications of Multilayer Protocols . . . 193
Converged Protocols . . . 194
Micro-segmentation . . . 195
Wireless Technologies . . . 197
Wireless Theory and Signaling . . . 197
Wi-Fi . . . 199
Bluetooth . . . 202
Zigbee . . . 202
Satellite . . . 203
Li-Fi . . . 203
Cellular Networks . . . 204
Content Distribution Networks . . . 205
REVIEW . . . 206
4.1 QUESTIONS . . . 206
4.1 ANSWERS . . . 207

Objective 4.2 Secure network components . . . 207
Network Security Design and Components . . . 208
Operation of Hardware . . . 208
Transmission Media . . . 212
Endpoint Security . . . 213
REVIEW . . . 214
4.2 QUESTIONS . . . 214
4.2 ANSWERS . . . 214

Objective 4.3 Implement secure communication channels according to design . . . 215
Securing Communications Channels . . . 215
Voice . . . 215
Multimedia Collaboration . . . 218

Remote Access 219
Data Communications 220
Virtualized Networks 222
Third-Party Connectivity 222
REVIEW 223
4.3 QUESTIONS 223
4.3 ANSWERS 224

5.0 Identity and Access Management (IAM) 225

Objective 5.1 Control physical and logical access to assets 226
Controlling Logical and Physical Access 226
Logical Access 227
Physical Access 228
REVIEW 228
5.1 QUESTIONS 228
5.1 ANSWERS 229

Objective 5.2 Manage identification and authentication of people, devices, and services 229
Identification and Authentication 229
Identity Management Implementation 230
Single/Multifactor Authentication 230
Accountability 231
Session Management 232
Registration, Proofing, and Establishment of Identity 232
Federated Identity Management 233
Credential Management Systems 233
Single Sign-On 234
Just-in-Time 234
REVIEW 235
5.2 QUESTIONS 236
5.2 ANSWERS 236

Objective 5.3 Federated identity with a third-party service 237
Third-Party Identity Services 237
On-Premise 237
Cloud 238
Hybrid 238
REVIEW 238
5.3 QUESTIONS 239
5.3 ANSWERS 239

Objective 5.4 Implement and manage authorization mechanisms. . . 239
Authorization Mechanisms and Models . . . 240
Discretionary Access Control . . . 241
Mandatory Access Control . . . 241
Role-Based Access Control . . . 242
Rule-Based Access Control . . . 242
Attribute-Based Access Control . . . 243
Risk-Based Access Control . . . 243
REVIEW . . . 243
5.4 QUESTIONS . . . 244
5.4 ANSWERS . . . 244
Objective 5.5 Manage the identity and access provisioning lifecycle . . . 245
Identity and Access Provisioning Life Cycle . . . 245
Provisioning and Deprovisioning . . . 245
Role Definition . . . 247
Privilege Escalation . . . 248
Account Access Review . . . 249
REVIEW . . . 251
5.5 QUESTIONS . . . 251
5.5 ANSWERS . . . 252
Objective 5.6 Implement authentication systems . . . 252
Authentication Systems . . . 252
Open Authorization . . . 253
OpenID Connect . . . 253
Security Assertion Markup Language . . . 253
Kerberos . . . 254
Remote Access Authentication and Authorization . . . 256
REVIEW . . . 257
5.6 QUESTIONS . . . 257
5.6 ANSWERS . . . 258
6.0 Security Assessment and Testing . . . 259
Objective 6.1 Design and validate assessment, test, and audit strategies . . . 260
Defining Assessments, Tests, and Audits . . . 260
Designing and Validating Evaluations . . . 261
Goals and Strategies . . . 261
Use of Internal, External, and Third-Party Assessors . . . 262
REVIEW . . . 263
6.1 QUESTIONS . . . 263
6.1 ANSWERS . . . 264

Objective 6.2 Conduct security control testing 264
Security Control Testing 264
Vulnerability Assessment 265
Penetration Testing 265
Log Reviews 267
Synthetic Transactions 268
Code Review and Testing 268
Misuse Case Testing 269
Test Coverage Analysis 269
Interface Testing 269
Breach Attack Simulations 270
Compliance Checks 270
REVIEW 271
6.2 QUESTIONS 271
6.2 ANSWERS 272
Objective 6.3 Collect security process data (e.g., technical and administrative) 272
Security Data 272
Security Process Data 273
REVIEW 275
6.3 QUESTIONS 276
6.3 ANSWERS 276
Objective 6.4 Analyze test output and generate report 277
Test Results and Reporting 277
Analyzing the Test Results 277
Reporting 278
Remediation, Exception Handling, and Ethical Disclosure 278
REVIEW 280
6.4 QUESTIONS 280
6.4 ANSWERS 280
Objective 6.5 Conduct or facilitate security audits 281
Conducting Security Audits 281
Internal Security Auditors 282
External Security Auditors 282
Third-Party Security Auditors 283
REVIEW 284
6.5 QUESTIONS 284
6.5 ANSWERS 284

7.0 Security Operations 285

Objective 7.1 Understand and comply with investigations 286
Investigations 286
Forensic Investigations 287
Evidence Collection and Handling 287
Digital Forensics Tools, Tactics, and Procedures 290
Investigative Techniques 291
Reporting and Documentation 292
REVIEW 293
7.1 QUESTIONS 294
7.1 ANSWERS 294

Objective 7.2 Conduct logging and monitoring activities 295
Logging and Monitoring 295
Continuous Monitoring 296
Intrusion Detection and Prevention 296
Security Information and Event Management 297
Egress Monitoring 297
Log Management 298
Threat Intelligence 298
User and Entity Behavior Analytics 301
REVIEW 302
7.2 QUESTIONS 303
7.2 ANSWERS 304

Objective 7.3 Perform Configuration Management (CM) (e.g., provisioning, baselining, automation) 304
Configuration Management Activities 304
Provisioning 305
Baselining 305
Automating the Configuration Management Process 306
REVIEW 306
7.3 QUESTIONS 307
7.3 ANSWERS 307

Objective 7.4 Apply foundational security operations concepts 308
Security Operations 308
Need-to-Know/Least Privilege 308
Separation of Duties and Responsibilities 309
Privileged Account Management 310
Job Rotation 311
Service Level Agreements 312
REVIEW 313
7.4 QUESTIONS 314
7.4 ANSWERS 314

Objective 7.5 Apply resource protection 314
Media Management and Protection 315
Media Management 315
Media Protection Techniques 315
REVIEW 317
7.5 QUESTIONS 317
7.5 ANSWERS 318
Objective 7.6 Conduct incident management 318
Security Incident Management 318
Incident Management Life Cycle 319
REVIEW 324
7.6 QUESTIONS 325
7.6 ANSWERS 326
Objective 7.7 Operate and maintain detective and preventative measures 326
Detective and Preventive Controls 326
Allow-Listing and Deny-Listing 327
Firewalls 328
Intrusion Detection Systems and Intrusion Prevention Systems 331
Third-Party Provided Security Services 332
Honeypots and Honeynets 333
Anti-malware 334
Sandboxing 335
Machine Learning and Artificial Intelligence 336
REVIEW 336
7.7 QUESTIONS 338
7.7 ANSWERS 338
Objective 7.8 Implement and support patch and vulnerability management 338
Patch and Vulnerability Management 339
Managing Vulnerabilities 339
Managing Patches and Updates 340
REVIEW 342
7.8 QUESTIONS 342
7.8 ANSWERS 343
Objective 7.9 Understand and participate in change management processes 344
Change Management 344
Change Management Processes 344
REVIEW 347
7.9 QUESTIONS 347
7.9 ANSWERS 348

Objective 7.10 Implement recovery strategies 348
Recovery Strategies 348
Backup Storage Strategies 348
Recovery Site Strategies 351
Multiple Processing Sites 352
Resiliency 355
High Availability 355
Quality of Service 356
Fault Tolerance 356
REVIEW 357
7.10 QUESTIONS 358
7.10 ANSWERS 359
Objective 7.11 Implement Disaster Recovery (DR) processes. 359
Disaster Recovery 359
Saving Lives and Preventing Harm to People 360
The Disaster Recovery Plan 360
Response 361
Personnel 361
Communications 361
Assessment 363
Restoration 363
Training and Awareness 364
Lessons Learned 364
REVIEW 365
7.11 QUESTIONS 366
7.11 ANSWERS 367
Objective 7.12 Test Disaster Recovery Plans (DRP). 367
Testing the Disaster Recovery Plan 367
Read-Through/Tabletop 368
Walk-Through 369
Simulation 369
Parallel Testing 370
Full Interruption 370
REVIEW 371
7.12 QUESTIONS 371
7.12 ANSWERS 372
Objective 7.13 Participate in Business Continuity (BC) planning and exercises 372
Business Continuity 372
Business Continuity Planning 373
Business Continuity Exercises 375

REVIEW 376
7.13 QUESTIONS 376
7.13 ANSWERS 377
Objective 7.14 Implement and manage physical security 377
Physical Security 377
Perimeter Security Controls 378
Internal Security Controls 382
REVIEW 386
7.14 QUESTIONS 387
7.14 ANSWERS 387
Objective 7.15 Address personnel safety and security concerns 388
Personnel Safety and Security 388
Travel 388
Security Training and Awareness 389
Emergency Management 389
Duress 390
REVIEW 391
7.15 QUESTIONS 391
7.15 ANSWERS 392
8.0 Software Development Security 393
Objective 8.1 Understand and integrate security in the Software Development Life Cycle (SDLC) 394
Software Development Life Cycle 394
Development Methodologies 395
Maturity Models 398
Operation and Maintenance 400
Change Management 401
Integrated Product Team 401
REVIEW 401
8.1 QUESTIONS 402
8.1 ANSWERS 403
Objective 8.2 Identify and apply security controls in software development ecosystems 403
Security Controls in Software Development 403
Programming Languages 404
Libraries 405
Tool Sets 406
Integrated Development Environment 406
Runtime 406
Continuous Integration and Continuous Delivery 407
Security Orchestration, Automation, and Response 407
Software Configuration Management 408

Code Repositories 408
Application Security Testing 408
REVIEW 411
8.2 QUESTIONS 411
8.2 ANSWERS 412
Objective 8.3 Assess the effectiveness of software security 412
Software Security Effectiveness 412
Auditing and Logging Changes 413
Risk Analysis and Mitigation 413
REVIEW 415
8.3 QUESTIONS 415
8.3 ANSWERS 415
Objective 8.4 Assess security impact of acquired software 416
Security Impact of Acquired Software 416
Commercial-off-the-Shelf Software 416
Open-Source Software 417
Third-Party Software 417
Managed Services 418
REVIEW 419
8.4 QUESTIONS 419
8.4 ANSWERS 420
Objective 8.5 Define and apply secure coding guidelines and standards 420
Secure Coding Guidelines and Standards 420
Security Weaknesses and Vulnerabilities at the Source-Code Level 420
Security of Application Programming Interfaces 421
Secure Coding Practices 422
Software-Defined Security 424
REVIEW 424
8.5 QUESTIONS 425
8.5 ANSWERS 425

A About the Online Content 427

System Requirements 427
Your Total Seminars Training Hub Account 427
Privacy Notice 427
Single User License Terms and Conditions 427
TotalTester Online 429
Technical Support 429

Index 431

This page intentionally left blank

Acknowledgments

A book isn't simply written by one person; so many people had key roles in the production of this study guide, so I'd like to take this opportunity to acknowledge and thank them. First and foremost, I would like to thank the folks at McGraw Hill, Wendy Rinaldi, Caitlin Cromley-Linn, and Janet Walden. All three worked hard to keep me on track and made sure that this book met the highest standards of quality. They are awesome people to work with, and I'm grateful once again to work with them!

I would also like to sincerely thank Nitesh Sharma, Senior Project Manager, Knowledge-Works Global Ltd, who worked on the post-production for the book, and Bill McManus, who did the copyediting work for the book. They are also great folks to work with. Nitesh was so patient and professional with me at various times when I did not exactly meet a deadline and I'm so grateful for that. I've worked with Bill a few times on different book projects, and I must admit I'm always in awe of him (and a bit intimidated by him, but really glad in the end to have him help on my projects), since he is an awesome copyeditor who catches every single one of the plentiful mistakes I make during the writing process. I have also gained a significant respect for Bill's knowledge of cybersecurity, as he's always been able to key in on small nuances of wonky explanations that even I didn't catch and suggest better ways to write them. He's the perfect person to make sure this book flows well, is understandable to a reader, and is a higher-quality resource. Thank you, Bill!

There are many other people on the production side who contributed significantly to the publication of this book, including Rachel Fogelberg, Ted Laux, Thomas Somers, and Jeff Weeks, as well as others. My sincere thanks to them all for their hard work.

I also want to thank my family for their patience and understanding as I took time away from them to write this book. I owe them a great deal of time I can never pay back, and I am very grateful for their love and support.

And last, but certainly not least, I want to thank the technical editor, Nichole O'Brien. I've worked with Nichole on tons of real-world cybersecurity projects off and on for at least ten years now. I've lost count of how many proposals, risk assessment reports, customer meetings, and cyber-related problems she has suffered through with me, yet she didn't hesitate to jump in and become the technical editor for this book. Nichole is absolutely one of the smartest businesspeople I know in cybersecurity, as well as simply a really good person, and I have an infinite amount of professional and personal respect for her. This book is so much better for having her there to correct my mistakes, ask critical questions, make me do more research, and add a different and unique perspective to the process. Thanks, Nichole!

—Bobby Rogers

Introduction

Welcome to *CISSP Passport*! This book is focused on helping you to pass the Certified Information Systems Security Professional (CISSP) certification examination from the International Information System Security Certification Consortium, or (ISC)². The idea behind the Passport series is to give you a concise study guide for learning the key elements of the certification exam from the perspective of the required objectives published by (ISC)², in their CISSP Certification Exam Outline. Cybersecurity professionals can review the experience requirements set forth by (ISC)² at https://www.isc2.org/Certifications/CISSP/experience-requirements. The basic requirement is five years of cumulative paid work experience in two or more of the eight CISSP domains, or four years of such experience plus either a four-year college degree or an additional credential from the (ISC)² approved list. (ISC)² requires that you document this experience before you can be fully certified as a CISSP. For those candidates who do not yet meet the experience requirements, they may achieve Associate of (ISC)² status by passing the examination. Associates of (ISC)² are then allowed up to six years to accumulate the required five years of experience to become full CISSPs.

The eight domains and the approximate percentage of exam questions they represent are as follows:

- Security and Risk Management (15%)
- Asset Security (10%)
- Security Architecture and Engineering (13%)
- Communication and Network Security (13%)
- Identity and Access Management (IAM) (13%)
- Security Assessment and Testing (12%)
- Security Operations (13%)
- Software Development Security (11%)

CISSP Passport assumes that you have already studied long and hard for the CISSP exam and now just need a quick refresher before you take the exam. This book is meant to be a "no fluff" concise study guide with quick facts, definitions, memory aids, charts, and brief explanations. Because this guide gives you the key concepts and facts, and not the in-depth

explanations surrounding those facts, you should *not* use this guide as your only study source to prepare for the CISSP exam. There are numerous books you can use for your deep studying, such as *CISSP All-in-One Exam Guide, Ninth Edition*, also from McGraw Hill.

I recommend that you use this guide to reinforce your knowledge of key terms and concepts and to review the broad scope of topics quickly in the final few days before your CISSP exam, after you've done all of your "deep" studying. This guide will help you memorize fast facts, as well as refresh your memory about topics you may not have studied for a while.

This guide is organized around the most recent CISSP exam domains and objectives released by (ISC)², which is May 1, 2021 at the time of writing this book. Keep in mind that (ISC)² reserves the right to change or update the exam objectives anytime at its sole discretion and without any prior notice, so you should check the (ISC)² website for any recent changes before you begin reading this guide and again a week or so before taking the exam to make sure you are studying the most updated materials.

The structure of this study guide parallels the structure of the eight CISSP domains published by (ISC)², presented in the same numerical order in the book, with individual domain objectives also ordered by objective number in each domain. Each domain in this guide is equivalent to a regular book chapter, so this guide has eight considerably large "chapters" with individual sections devoted to the objective numbers. This organization is intended to help you learn and master each objective in a logical way. Because some domain objectives overlap, you will see a bit of redundancy in topics discussed throughout the book; where this is the case, the topic is presented in its proper context within the current domain objective and you'll see a cross-reference to the other objective(s) in which the same topic is discussed.

Each domain contains the following useful items to call out points of interest.



EXAM TIP Indicates critical topics you're likely to see on the actual exam.



NOTE Points out ancillary but pertinent information, as well as areas for further study.



CAUTION Warns you of common pitfalls, misconceptions, and potentially harmful or risky situations when working with the technology in the real world.

Cross-Reference

Directs you to other places in the book where concepts are covered, for your reference.



ADDITIONAL RESOURCES Identifies where you can find books, websites, and other media for further assistance.

The end of each objective gives you two handy tools. The "Review" section provides a synopsis of the objective—a great way to quickly review the critical information. Then the "Questions" and "Answers" sections enable you to test your newly acquired knowledge. For further study, this book includes access to online practice exams that will help to prepare you for taking the exam itself. All the information you need for accessing the exam questions is provided in the appendix. I recommend that you take the practice exams to identify where you have knowledge gaps and then go back and review the relevant material as needed.

I hope this book is helpful to you not only in studying for the CISSP exam but also as a quick reference guide you'll use in your professional life. Thanks for picking this book to help you study, and good luck on the exam!

This page intentionally left blank



Security and Risk Management

Domain Objectives

- **1.1** Understand, adhere to, and promote professional ethics.
- **1.2** Understand and apply security concepts.
- **1.3** Evaluate and apply security governance principles.
- **1.4** Determine compliance and other requirements.
- **1.5** Understand legal and regulatory issues that pertain to information security in a holistic context.
- **1.6** Understand requirements for investigation types (i.e., administrative, criminal, civil, regulatory, industry standards).
- **1.7** Develop, document, and implement security policy, standards, procedures, and guidelines.
- **1.8** Identify, analyze, and prioritize Business Continuity (BC) requirements.
- **1.9** Contribute to and enforce personnel security policies and procedures.
- **1.10** Understand and apply risk management concepts.
- **1.11** Understand and apply threat modeling concepts and methodologies.
- **1.12** Apply Supply Chain Risk Management (SCRM) concepts.
- **1.13** Establish and maintain a security awareness, education, and training program.

Domain 1, "Security and Risk Management," is one of the key domains in understanding critical security principles that you will encounter on the CISSP exam. The majority of the topics in this domain include the administrative or managerial security measures put in place to manage a security program. In this domain you will learn about professional ethics and important fundamental security concepts. We will discuss governance and compliance, investigations, security policies, and other critical management concepts. We will also delve into business continuity, personnel security, and the all-important risk management processes. We'll also discuss threat modeling, explore supply chain risk management, and finish the domain by examining the different aspects of security training and awareness programs. These are all very important concepts that will help you to understand the subsequent domains, since they provide the foundations of knowledge you need to be successful on the exam.



Objective 1.1 Understand, adhere to, and promote professional ethics

The fact that (ISC)² places professional ethics as the first objective in the first domain of the CISSP exam requirements speaks volumes about the importance of ethics and ethical behavior in our profession. The continuing increases in network breaches, data loss, and ransomware demonstrate the criticality of ethical conduct in this expanding information security landscape. Our information systems security workforce is expanding at a rapid pace, and these new recruits need to understand the professional discipline required to succeed. Some may enter the field because they expect to make a lot of money, but ultimately competence, integrity, and trustworthiness are the qualities necessary for success. Most professions have published standards for ethical behavior, such as healthcare, law enforcement, accounting, and many other professions. In fact, you would be hard-pressed to find a profession that does not have at least some type of minimal ethical requirements for professional conduct.

While exam objective 1.1 is the only objective that explicitly covers ethics and professional conduct, it's important to emphasize them, since you will be expected to know them on the exam and, more importantly, you will be expected to uphold them to maintain your CISSP status. The first part of this exam objective covers the core ethical requirements from (ISC)² itself. Absent any other ethical standards that you may also be required to uphold in your profession, from your organization, your customers, and even any other certifications you hold, the (ISC)² Code of Ethics should be sufficient to guide you in ethical behavior and professional conduct while you are employed as an information systems security professional for as long as you hold the CISSP certification. The second part of the objective reviews other sources of professional ethics that guide your conduct, such as those from industry or professional organizations. First, let's look at the (ISC)² Code of Ethics.

The (ISC)² Code of Ethics

The (ISC)² Code of Ethics, located on the (ISC)² website at https://www.isc2.org/Ethics#, consists of a preamble and four mandatory canons. Additionally, the web page includes a comprehensive set of ethics complaint procedures for filing ethics complaints against certified members. The complaint procedures are designed to detail how someone might formally accuse a certified member of violating one or more of the four canons.



NOTE (ISC)² updates the Code of Ethics from time to time, so it is best to occasionally go to the (ISC)² website and review it for any changes. This allows you to keep up with current requirements and serves to remind you of your ethical and professional responsibilities.

Code of Ethics Preamble

The Code of Ethics Preamble simply states that people who are bound to the code must adhere to the highest ethical standards of behavior, and that the code is a condition of certification. Per the (ISC)² site (https://www.isc2.org/Ethics#), the preamble states (at the time of writing):

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

Code of Ethics Canons

The Code of Ethics Canons dictate the more specific requirements that certification holders must obey. According to the ethics complaint procedures detailed by (ISC)², violation of any of these canons is grounds for the certificate holder have their certification revoked. The canons are as follows:

I. Protect society, the common good, necessary public trust and confidence, and the infrastructure.

II. Act honorably, honestly, justly, responsibly, and legally.

III. Provide diligent and competent service to principals.

IV. Advance and protect the profession.

Obviously, these canons are intentionally broad and, unfortunately, someone could construe them to fit almost any type of act by a CISSP, accidental or malicious, into one these categories. However, the ethics complaint procedures specify a burden of proof involved with making a complaint against the certification holder for violation of these canons. The complaint procedures, set forth in the "Standing of Complainant" section, specify that "complaints

will be accepted only from those who claim to be injured by the alleged behavior." Anyone with knowledge of a breach of Canons I or II may file a complaint against someone, but only principals, which are employers or customers of the certificate holder, can lodge a complaint about any violation of Canon III, and only other certified professionals may register complaints about violations of Canon IV.

Also according to the ethics complaint procedures, the complaint goes before an ethics committee, which hears complaints of breaches of the Code of Ethics Canons, and makes a recommendation to the board. But the board ultimately makes decisions regarding the validity of complaints, as well as levees the final disciplinary action against the member, if warranted. A person who has had an ethics complaint lodged against them under these four canons has a right to respond and comment on the allegations, as there are sound due process procedures built into this process.



EXAM TIP You should be familiar with the preamble and the four canons of the (ISC)² Code of Ethics for the exam. It's a good idea to go to the (ISC)² website and review the most current Code of Ethics shortly before you take the exam.

Organizational Code of Ethics

The second part of exam objective 1.1 encompasses organizational standards and codes of ethics. Most organizations today have some minimal form of a code of ethics, professional standards, or behavioral requirements that you must obey to be a member of that organization. "Organization" in this context means professional organizations, your workplace, your customer organization, or any other formal, organized body to which you belong or are employed by. Whether you are a government employee or a private contractor, whether you work for a volunteer agency or work in a commercial setting, you're likely required to adhere to some type of organizational code of ethics. Let's examine some of the core requirements most organizational codes of ethics have in common.

Workplace Ethics Statements and Policies

Codes of ethics in the workplace may or may not be documented. Often there is no formalized, explicit code of ethics document published by the organization, although that may not be the case, especially in large or publicly traded corporations. More often than not, the requirements for ethical or professional behavior are stated as a policy or group of policies that apply not only to the security professionals in the organization but to every employee. For example, there are usually policies that cover the topics of acceptable use of organizational IT assets, personal behavior toward others, sexual harassment and bullying, bribery, gifts from external parties, and so on. Combined, these policies cover the wide range of professional behavior expectations. These policies may be sponsored and monitored by the human resources department and are likely found in the organization's employee handbook. For the organizations that have

explicit professional ethics documents, these usually describe general statements that are not specific to IT or cybersecurity professionals and direct the employee to behave ethically and professionally in all matters.

Other Sources for Ethics Requirements

Although not directly testable by the CISSP exam, it's worth noting that there are other sources for ethics requirements for technology professionals in general and cybersecurity professionals in particular. All of these sources contain similar requirements to act in a professional, honest manner while protecting the interests of customers, employers, and other stakeholders, as well as maintain professional integrity and work toward the good of society. The following subsections describe several sources of professional ethics standards to give you an idea of how important ethics and professional behavior are across the wide spectrum of not only cybersecurity but technology in general.

The Computer Ethics Institute

The Computer Ethics Institute (CEI) is a nonprofit policy, education, and research group founded to promote the study of technology ethics. Its membership includes several technology-related organizations and prominent technologists and it is positioned as a forum for public discussion on a variety of topics affecting the integration of technology and society. The most well-known of its efforts is the development of the Ten Commandments of Computer Ethics, which has been used as the basis of several professional codes of ethics and behavior documents, among them the (ISC)² Code of Ethics.

The Ten Commandments of Computer Ethics, presented here from the CEI website, are as follows:

1. Thou shalt not use a computer to harm other people.
2. Thou shalt not interfere with other people's computer work.
3. Thou shalt not snoop around in other people's computer files.
4. Thou shalt not use a computer to steal.
5. Thou shalt not use a computer to bear false witness.
6. Thou shalt not copy or use proprietary software for which you have not paid.
7. Thou shalt not use other people's computer resources without authorization or proper compensation.
8. Thou shalt not appropriate other people's intellectual output.
9. Thou shalt think about the social consequences of the program you are writing or the system you are designing.
10. Thou shalt always use a computer in ways that ensure consideration and respect for your fellow humans.

Institute of Electrical and Electronics Engineers – Computer Society

The Institute of Electrical and Electronics Engineers (IEEE) published a professional Code of Ethics designed to promulgate ethical behaviors among technology professionals. Although the IEEE Code of Ethics does not specifically target cybersecurity professionals, its principles similarly promote the professional and ethical behaviors of other technology professionals and is similar in requirements to the (ISC)² Code of Ethics. The more important points of the IEEE Code of Ethics are summarized as follows:

- Uphold high standards of integrity, responsible behavior, and ethical conduct in professional activities
- Hold paramount the safety, health, and welfare of the public
- Avoid real or perceived conflicts of interest
- Avoid unlawful conduct
- Treat all persons fairly and with respect
- Ensure the code is upheld by colleagues and coworkers

As you can see, these points are directly aligned with the (ISC)² Code of Ethics and, as with many codes of conduct, offer no conflict with other codes that members may be subject to. In fact, since codes of ethics and professional behavior are often similar, they support and serve to strengthen the requirements levied on various individuals.



ADDITIONAL RESOURCES In addition to the example of the IEEE Code of Ethics, numerous other professional organizations that are closely related to or aligned with cybersecurity professionals also have comparable codes that are worth mentioning. Another noteworthy example is the Project Management Institute (PMI) Code of Ethics and Professional Conduct, available at https://www.pmi.org/about/ethics/code.

Governance Ethics Requirements

There also are standards that are imposed as part of regulatory requirements that cover how technology professionals will comport themselves. Some of these standards don't specifically target cybersecurity professionals per se, but they do prescribe ethical behaviors with regard to data protection, for example, and apply to organizations and personnel alike. Almost all data protection regulations, such as the EU's General Data Protection Regulation (GDPR), the U.S. Health Insurance Portability and Accountability Act (HIPAA), the National Institute of Standards and Technology (NIST) publications, the Code of Ethics requirements spelled out in Section 406 of the Sarbanes-Oxley Act of 2002, and countless other laws and regulations, describe the actions that users and personnel with privileged access to sensitive data must take to protect that data from a legal and ethical perspective in order to comply with security, privacy, and other governance requirements.

REVIEW

Objective 1.1: Understand, adhere to, and promote professional ethics In this objective we focused on one of the more important objectives for the CISSP exam—one that's often overlooked in exam prep. We discussed codes of ethics, which are requirements intended to guide our professional behavior. We specifically examined the (ISC)² Code of Ethics, as that is the most relevant to the exam. The Code of Ethics consists of a preamble and four mandatory canons. (ISC)² also has a comprehensive set of complaint procedures for ethics complaints against certified members. The complaint procedures detail the process for formally accusing a certified member of violating one or more of the four canons, while ensuring a fair and impartial due process for the accused.

We also examined organizational ethics and discussed how some organizations may not have a formalized code of ethics document, but their ethical or professional behavior expectations may be contained in their policies. These are usually found in policies such as acceptable use, acceptance of gifts, bribery, and other types of policies. Most of the policies that affect professional behavior for employees are typically found in the employee handbook.

Finally, we discussed other sources of professional ethics, from professional organizations and governance requirements that may define how to protect certain sensitive data classifications. Absent any other core ethics document that prescribes professional behavior, the (ISC)² Code of Ethics is mandatory for CISSP certification holders and should be used to guide their behavior.

1.1 QUESTIONS

1. You're a CISSP who works for a small business. Your workplace has no formalized code of professional ethics. Your manager recently asked you to fudge the results of a vulnerability assessment on a group of production servers to make it appear as if the security posture is improving. Absent a workplace code of ethics, which of the following should guide your behavior regarding this request?
   - **A.** Your own professional conscience
   - **B.** (ISC)² Code of Ethics
   - **C.** Workplace Acceptable Use Policy
   - **D.** The Computer Ethics Institute policies
2. Nichole is a security operations center (SOC) supervisor who has observed one of her CISSP-certified subordinates in repeated violation of both the company's requirements for professional behavior and the (ISC)² Code of Ethics. Which of the following actions should she take?
   - **A.** Report the violation to the company's HR department only
   - **B.** Report the violation to (ISC)² and the HR department
   - **C.** Ignore a one-time violation and counsel the individual
   - **D.** Report the violation to (ISC)² only

3. Which of the following is a legal, ethical, or professional requirement levied upon an individual to protect data based upon the specific industry, data type, and sensitivity?
   - **A.** (ISC)² Code of Ethics
   - **B.** IEEE Code of Ethics
   - **C.** The Sarbanes-Oxley Code of Ethics requirements
   - **D.** The Computer Ethics Institute's Ten Commandments of Computer Ethics
4. Bobby has been accused of violating one of the four canons of the (ISC)² Code of Ethics. A fellow cybersecurity professional has made the complaint that Bobby intentionally wrote a cybersecurity audit report to reflect favorably on a company in which he is also applying for a job in order to gain favor with its managers. Which of the following four canons has Bobby likely violated?
   - **A.** Provide diligent and competent service to principals
   - **B.** Act honorably, honestly, justly, responsibly, and legally
   - **C.** Advance and protect the profession
   - **D.** Protect society, the common good, necessary public trust and confidence, and the infrastructure

1.1 ANSWERS

1. **B** Absent any other binding code of professional ethics from the workplace, the (ISC)² Code of Ethics binds certified professionals to a higher standard of behavior. While using your own professional judgment is admirable, not everyone's professional standards are at the same level. Workplace policies do not always cover professional conduct by cybersecurity personnel specifically. The Computer Ethics Institute policies are not binding to cybersecurity professionals.
2. **B** Since the employee has violated both the company's professional behavior requirements and the (ISC)² Code of Ethics, Nichole should report the actions to both entities. Had the violation been only that of the (ISC)² Code of Ethics, she would not have necessarily needed to report it to the company. One-time violations may be accidental and should be handled at the supervisor's discretion; however, repeated violations may warrant further action depending upon the nature of the violation and the situation.
3. **C** The Sarbanes-Oxley (SOX) Code of Ethics requirements are part of the regulation (Section 406 of the Act) enacted to prevent securities and financial fraud and require organizations to enact codes of ethics to protect financial and personal data. The other choices are not focused on data sensitivity or regulations, but rather apply to technology and cybersecurity professionals.
4. **A** Although the argument can be made that falsifying an audit report could violate any or all of the four (ISC)² Code of Ethics Canons, the scenario specifically affects the canon that requires professionals to perform diligent and competent service to principals.



Objective 1.2 Understand and apply security concepts

In this objective we will examine some of the more fundamental concepts of security. Although fundamental, they are critical in understanding everything that follows, since everything we will discuss in future objectives throughout all CISSP domains relates to the goals of security and their supporting tenets.

Security Concepts

To become certified as a CISSP, you must have knowledge and experience that covers a wide variety of topics. However, regardless of the experience you may have in the different domains, such as networking, digital forensics, compliance, or penetration testing, you need to comprehend some fundamental concepts that are the basis of all the other security knowledge you will need in your career. This core knowledge includes the goals of security and its supporting principles.

In this objective we're going to discuss this core knowledge, which serves as a reminder for the experience you likely already have before attempting the exam. We'll cover the goals of security as well as the supporting tenets, such as identification, authentication, authorization, and nonrepudiation. We will also discuss key supporting concepts such as principles of least privilege and separation of duties. You'll find that no matter what expertise you have in the CISSP domains, these core principles are the basis for all of them. As we discuss each of these core subjects we'll talk about how different topics within the CISSP domains articulate to these areas. First, it's useful to establish common ground with some terms you'll likely see throughout this book and your studies for the exam.

Data, Information, Systems, and Entities

There are terms that we commonly use in cybersecurity that can cause confusion if everyone in the field does not have a mutual understanding of what the terms mean. Our field is rich with acronyms, such as MAC, DAC, RBAC, IdM, and many more. Often the same acronym can stand for different terms. For example, in information technology and cybersecurity parlance, MAC can stand for media access control, message authentication code, mandatory access control, and memory access controller, not to mention that it's also a slang term for a Macintosh computer. That's an example of why it's important to define a few terms up front before we get into our discussion of security concepts. These terms include data, information, system, and entity (and its related terms subject and object).

Two terms often used interchangeably by technology people in everyday conversation are *data* and *information*. In nontechnical discussion, the difference really doesn't matter, but as cybersecurity professionals, we need to be more precise in our speech and differentiate

between the two. For purposes of this book, and studying for the exam, *data* are raw, singular pieces of fact or knowledge that have no immediate context or meaning. An example might be an IP address, or domain name, or even an audit log entry, which by itself may not have any meaning. *Information* is data organized into context and given meaning. An example might be several pieces of data that are correlated to show an event that occurred on host at a specific time by a specific individual.



EXAM TIP The CISSP exam objectives do not distinguish the differences between the terms "information" and "data," as they are often used interchangeably in the profession as well. For the purposes of this book, we also will sometimes not distinguish the difference and use the term interchangeably, depending on the context and the exam objectives presented.

A *system* consists of multiple components such as hardware, software, network protocols, and even processes. A system could also consist of multiple smaller systems, sometimes called a system of systems but most frequently just referred to as a system, regardless of the type or quantity of subsystems.

An *entity*, for our purposes, is a general, abstract term that includes any combination of organizations, persons, hardware, software, processes, and so on, that may interact with people, systems, information, or data. Frequently we talk about users accessing data, but in reality, software programs, hardware, and processes can also independently access data and other resources on a network, regardless of user action. So it's probably more correct to say that an entity or entities access these resources. We can assign accounts and permissions to almost any type of entity, not just humans. It's also worth noting that entities are also referred to as *subjects*, which perform actions (read, write, create, delete, etc.) on *objects*, which are resources such as computers, systems, and information.

Now that we have those terms defined, let's discuss the three goals of security—confidentiality, integrity, and availability.

Confidentiality

Of the three primary goals of information security, confidentiality is likely the one that most people associate with cybersecurity. Certainly, it's important to make sure that systems and data are kept confidential and only accessed by entities that have a valid reason, but the other goals of security, which we will discuss shortly, are also of equal importance. *Confidentiality* is about keeping information secret and, in some cases, private. It requires protecting information that is not generally accessible to everyone, but rather only to a select few. Whether it's personal privacy or health data, proprietary company information, classified government data, or just simply data of a sensitive nature, confidential information is meant to be kept secret. In later objectives we will discuss different access controls, such as file permissions, encryption, authentication schemes, and other measures, that are designed to keep data and systems confidential.

Integrity

Integrity is the goal of security to ensure that data and systems are not modified or destroyed without authorization. To maintain integrity, data should be altered only by an entity that has the appropriate access and a valid reason to modify. Obviously, data may be altered purposefully for malicious reasons, but accidental or unintentional changes may be caused by a well-intentioned user or even by a bad network connection that degrades the integrity of a file or data transmission. Integrity is assured through several means, including identification and authentication mechanisms (discussed shortly), cryptographic methods (e.g., file hashing), and checksums.

Availability

Availability means having information and the systems that process it readily accessible by authorized users any time and in any manner they require. Systems and information do users little good if they can't get to and use those resources when needed, and simply preventing their authorized use contradicts the availability goal. Availability can be denied accidentally by a network or device outage, or intentionally by a malicious entity that destroys systems and data or prevents use via denial-of-service attacks. Availability can be ensured through various means including equipment redundancy, data backups, access control, and so on.

Supporting Tenets of Information Security

Security tenets are processes that support the three goals of security. The security tenets are identification, authentication, authorization, auditing, accountability, and nonrepudiation. Note that these may be listed differently or include other principles, depending on the source of knowledge or the organization.

Identification

Identification is the act of presenting credentials that state (assert) the identity of an individual or entity. A *credential* is a piece of information (physical or electronic) that confirms the identity of the credential holder and is issued by an authoritative source. Examples of credentials used to identify an entity include a driver's license, passport, username and password combination, smart card, and so forth.

Authentication

Authentication occurs after identification and is the process of verifying that the credential presented matches the actual identity of the entity presenting it. Authentication typically occurs when an entity presents an identification and credential, and the system or network verifies that credential against a database of known identities and characteristics. If the identity and credential asserted matches an entry in the database, the entity is authenticated.

Once this occurs, an entity is considered authenticated to the system, but that does not mean that they have the ability to perform any actions with any resources. This is where the next step, authorization, comes in.

Authenticity

Authenticity goes hand-in-hand with authentication, in that it is the validation of a user, an action, a document, or other entity through verified means. User authenticity is established with strong authentication mechanisms, for example; an action's authenticity is established through auditing and accountability mechanisms, and a document's authenticity might be established through integrity checks such as hashing.

Authorization

Authorization occurs only after an entity has been authenticated. *Authorization* determines what actions the entity can take with a given resource, such as a computer, application, or network. Note that it is possible for an entity to be authenticated but have no authorization to take any action with a resource. Authorization is typically determined by considering an individual's job position, clearance level, and need-to-know status for a particular resource. Authorization can be granted by a system administrator, a resource owner, or another entity in authority. Authorization is often implemented in the form of permissions, rights, and privileges used to interact with resources, such as systems and information.



EXAM TIP Remember that authorization consists of the actions an individual can perform, and is based on their job duties, security clearance, and need-to-know,

Auditing and Accountability

Accountability is the ability to trace and hold an entity responsible for any actions that entity has taken with a resource. Accountability is typically achieved through auditing. *Auditing* is the process of reviewing all interactions between an entity and an object to evaluate the effectiveness of security controls. An example is auditing access to a network folder and being able to conclusively determine that user Gary deleted a particular document in that folder. Auditing would rule out that another user performed this action on that resource. Most resources, such as computers, data, and information, can be audited for a variety of actions, such as access, creation, deletion, and so forth. The most frequent manifestation of auditing is through audit trails or logs, which are generated by the system or object being audited and record all actions that any user takes with that system or object.

Nonrepudiation

To hold entities, such as users, accountable for the actions they perform on objects, we must be able to conclusively connect their identity to an event. Auditing is useful for recording

interactions with systems and data to determine who is accountable for those actions. However, we also want to be able to ensure that we can have such fidelity in audit logs that the user or entity cannot later deny that they took the action. If we suspect that audit logs, for example, have been tampered with, altered, or even faked, then we can't conclusively hold someone accountable for their actions. *Nonrepudiation* is the inability of an entity to deny that it performed a particular action; in other words, through auditing and other means, it can be conclusively proven that an entity took a particular action and the entity cannot deny it. There are various methods used to ensure nonrepudiation, including audit log security, strong identification and authentication mechanisms, and strong auditing processes.

Supporting Security Concepts

Along with the three primary goals of security and their supporting security tenets, several key concepts form the foundation for good security. These are security principles that ensure individuals can perform only the actions they are allowed to perform, that there is a stated reason for having access to systems and information, and that individuals are not allowed to have excessive authority over systems and information. We will discuss three of these security concepts in the remainder of this objective and cover others later as we progress through the domain.

Principle of Least Privilege

One of the oldest and most basic concepts in information security is the principle of least privilege. Quite simply, the *principle of least privilege* states that entities should only have the minimum level of rights, permissions, privileges, and authority needed over systems and information to do their job, and no more. This limitation prevents them from being able to take actions that are beyond their authority or outside the scope of their duties. This concept applies to routine users, administrators, managers, and anyone who has any level of access to systems and information. The principle of least privilege applies to a variety of settings in information security, including the minimum permissions needed to access objects, the lowest level of rights over systems, and a limited ability to take actions that could affect systems and network resources. There are various ways that we assure the principle of least privilege, including minimal permissions and restricted accounts.

Need-to-Know

Need-to-know is a security concept that is related to the principle of least privilege. While the principle of least privilege means that users are explicitly assigned only the bare minimum of abilities to take action on system and information objects, *need-to-know* means that users should not have access to information or systems, regardless of assigned abilities, unless they need that access for their job. For example, if a person does not have the proper permissions to access a shared folder, need-to-know also implies that they should not be told the contents of what's in that folder, since it may be sensitive information. Only when a person has a demonstrated need-to-know for information, and received approval from their supervisory chain, should they be considered for additional rights or privileges to get access to systems and data.

Separation of Duties

Separation of duties is another key concept in information security, one that you will see implemented in various ways. Even when users have a valid need-to-know for information and properly assigned access for the minimum rights, permissions, and privileges to do their job, they should not have the ability to perform certain critical functions unless it is in conjunction with another person. The intent of *separation of duties* is to deny the user the ability to perform important functions unchecked, thereby requiring the oversight of someone else to help prevent disastrous results. If an individual is allowed to perform selected critical functions alone, another individual should be required to double-check for accuracy or completeness. This approach prevents a rogue user from doing serious damage to systems and information in an organization.



EXAM TIP While the principles of least privilege, need-to-know, and separation of duties are similar and complementary to each other, they are not synonymous. Understand the subtle differences between these terms.

REVIEW

Objective 1.2: Understand and apply security concepts In this objective we discussed key security concepts, which include the goals of security and supporting tenants and concepts. We discussed confidentiality, integrity, and availability, and how they are supported by different access controls. We also discussed tenets such as identification, authentication, authorization, accountability, auditing, and nonrepudiation. Finally, we talked about key concepts such as the principle of least privilege, need-to-know, and separation of duties.

1.2 QUESTIONS

1. Emilia is a new cybersecurity intern who works in a security operations center. During a mentoring session with her supervisor, she is asked about the differences between authentication and authorization. Which of the following is her best response?
   - **A.** Authorization validates identities, and authentication allows individuals access to resources.
   - **B.** Authentication allows individuals access to resources and is the same thing as authorization.
   - **C.** Authentication validates identities, and authorization allows individuals access to resources.
   - **D.** Authentication is the act of presenting a user identity to a system, and authorization validates that identity.

2. Evey is a cybersecurity analyst who works at a major research facility. Over time, the network administration staff has accumulated broad sets of privileges, and management now fears that one individual would be able to do significant damage to the network infrastructure if they have malicious intent. Evey is trying to sort out the different rights, permissions, and privileges that each network administrator has amassed. Which of the following concepts should she implement to ensure that a single person cannot perform a critical, potentially damaging function alone without it being detected or completed by another individual?
   - **A.** Separation of duties
   - **B.** Need-to-know
   - **C.** Principle of least privilege
   - **D.** Authorization
3. Ben is a member of his company's incident response team. Recently the company detected that several critical files in a sensitive data share have been subtly altered without anyone's knowledge. Which of the following was violated?
   - **A.** Nonrepudiation
   - **B.** Confidentiality
   - **C.** Availability
   - **D.** Integrity
4. Sam is a newly certified CISSP who has been tasked with reviewing audit logs for access to sensitive files. He has discovered that auditing is not configured properly, so it is difficult to trace the actions performed on an object to a unique individual and conclusively prove that the individual took the action. Which of the following is not possible because of the current audit configuration?
   - **A.** Authentication
   - **B.** Nonrepudiation
   - **C.** Authorization
   - **D.** Integrity

1.2 ANSWERS

1. **C** Authentication validates an identity when it is presented to the system, and authorization dictates which actions the user is allowed to perform on resources after they have been authenticated.
2. **A** Evey must implement separation of duties to ensure that network administrators can only perform critical functions in conjunction with another person. This would eliminate the ability of a single person to significantly damage the infrastructure in the event they have malicious intent, since it would require another individual to check their actions or complete a critical task.

3. **D** Unauthorized changes to critical files indicate that their integrity has changed.
4. **B** Without the ability to conclusively connect the actions performed on an object to a unique user identity, the user can deny (repudiate) that they took an action. This not only prevents accountability but also fails to ensure nonrepudiation.



Objective 1.3 Evaluate and apply security governance principles

In Objective 1.3 we will discuss security governance principles, which are the bedrock of the security program.

Security Governance

Security governance can best be described as requirements imposed on an organization by both internal and external entities that prescribe how the organization will protect its assets, to include systems and information. Security governance dictates how the organization will manage risk, be compliant with regulatory requirements, and operate its IT and cybersecurity programs. In this objective we will discuss both internal and external governance and how security functions align to business requirements. We'll also talk about how organizational processes are shaped by security governance and how in turn these same processes support that governance. We will briefly discuss the different roles and responsibilities involved in managing cybersecurity within the organization. We'll also go over the need for security control frameworks in managing organizational risk and protecting assets. Finally, we will explore the concepts of due care and due diligence and why they are critical in reducing risk and liability.

External Governance

External governance originates from sources outside the organization. The organization cannot control or ignore external governance requirements, as they stem from various sources including laws, regulations, and industry standards. External governance largely dictates how an organization protects certain classes of data, such as healthcare data (as mandated by HIPAA), financial data, and personal information. External governance also directs how an organization will interact with agencies outside of the organization, such as regulatory bodies, standards organizations, business partners, customers, competitors, and so on. External governance is typically mandatory and not subject to change or disregard by the organization.

Internal Governance

Internal governance stems from within the organization in the form of policies, procedures, adopted standards, and guidelines. For the purposes of this objective note that internal

requirements, which are typically articulated in the form of security policy, exist to support external governance. For example, if there is an external law or regulation imposed on the organization, internal policies are then written to state how that law or regulation will be followed and enforced within the organization. Policies and other internal governance impose mandatory standards of behavior on the organization and its members, as determined by senior management. The development and administration of internal governance must align with the organization's stated strategy, mission, goals, and objectives, which we will briefly discuss next.

Cross-Reference

We will discuss internal governance components in depth later in Objective 1.7.

Alignment of Security Functions to Business Requirements

Security doesn't exist for its own sake. It is actually an enabler for the business, though businesspeople who are not involved in security often argue that point. Security, through the overall information technology strategy or information security strategy, should complement and support the organization's overall strategy, mission, goals, and objectives. Security functions are those activities, tasks, and processes that are used to support those organizational requirements.

Business Strategy and Security Strategy

The *mission* of the organization is its stated purpose, the reason it is in business in the first place. For example, the mission of an automobile parts manufacturer is to supply parts to the larger automobile companies. An organization's mission is ongoing and seldom changes. *Goals* are what the organization wants to accomplish to further its mission. Most goals are formulated to cover three general time frames, referred to as long-term, near-term, and short-term goals, which correspond to strategic, operational, and tactical goals. *Strategic goals* are longer-term initiatives, typically three to five years, that address "big picture" issues and ideas to support the organization's mission. *Operational goals* cover shorter timeframes, usually one to three years (near-term) and focus on requirements necessary to maintain smooth and successful operations. *Tactical goals* refer to the day-to-day or short-term activities that accomplish routine tasks.

The organization should also have an information technology strategy and a supporting cybersecurity strategy. These two strategy documents (or one if they are combined) support the strategic and operational business goals. For example, if the organization's business strategy describes a path for how the company will expand into other countries over the next five years, the IT and cybersecurity strategy align with those goals and describe how the infrastructure must evolve over that timeframe to support the expansion effort. Both operational and tactical IT and cybersecurity activities in turn support their respective strategies by supporting the different organizational processes that exist to carry out the mission.

Organizational Processes

All cybersecurity activities must integrate with and support organizational processes, whether they are high- or low-level, strategic, or tactical processes. In turn, cybersecurity ramifications must be considered when these organizational processes are developed and implemented. For example, launching a major new product line is a business decision that must be supported by IT infrastructure expansion and changes, as well as by cybersecurity activities to keep those new systems secure and interoperable. Likewise the personnel responsible for launching the new product line must consider cybersecurity requirements as it is being designed and implemented. Senior executives often form security governance committees to evaluate and provide feedback on how security will affect and is affected by new or existing business processes, ventures, capabilities, and so on. There is also certainly some level of risk that inherently comes with new business processes and ventures, which the organization's senior management must address.

Many key organizational processes are closely coupled with security infrastructure. Although there are far too many processes to mention all of them here, the exam objectives call out some specific ones, particularly acquisitions and divestitures (along with the previously mentioned governance committees). These two processes involve acquiring another organization or, conversely, splitting an organization into different parts, sometimes into two completely new and independent organizations. Let's discuss each of these briefly.

An *acquisition* occurs when an organization buys or merges with another organization. This transaction is critical to the security infrastructure for both organizations in that the infrastructure of each is likely quite different, especially in terms of governance, data types and sensitivity, and how each organization manages security and risk. For this reason, during the acquisition process, the organization that is acquiring the other organization must perform its due diligence and due care (as discussed later in this objective) by researching the security posture and infrastructure of the other organization. The acquiring organization must identify and document key personnel, processes, and infrastructure components of the organization to be acquired. Most importantly, the acquiring organization must identify and document threats, vulnerabilities, and other elements of risk, since the organization is acquiring not only the new organization but also its risks.

The same principle also applies to divestitures. A *divestiture* is when an organization is splitting up into new, independent organizations, and when this happens, the division of data, personnel, and infrastructure between them must be carefully considered. Of course, these aren't the only things that are divided up among the new organizations—risk is also inherited by each of the individual organizations. Often it's the same risk, but sometimes it may be different depending upon the business processes and assets distributed to each new organization.

Organizational Roles and Responsibilities

Within each organization, people are appointed to various roles and have the responsibility of fulfilling different security functions—some at a higher strategic management level and some at a lower operational level. Most of the roles at the senior management level are legally

TABLE 1.3-1 Key Organizational Security Roles and Responsibilities

Role	Responsibility
Chief information officer (CIO)	Member of executive management responsible for all information technology in the organization.
Chief security officer (CSO)	Member of executive management responsible for all security operations in the organization.
Chief information security officer (CISO)	Member of executive management responsible for all information security aspects of the organization; may work for either the CIO or the CSO.
Chief privacy officer (CPO)	Responsible for ensuring customer, organization, and employee personal data is kept secure and used properly.
Data owner	Senior manager accountable and responsible for a particular classification of data; determines data sensitivity and establishes access control rules for that classification of data. Directs the use of security controls to protect data.
Data custodian	Responsible for day-to-day implementation of security controls used to protect data.
System owner	Senior manager accountable and responsible for a particular system which may process various classifications of data owned by different owners. Directs security controls used to protect systems.
System/security administrator	Responsible for day-to-day implementation of security controls used to protect systems.
Security auditor	Periodically checks to ensure that all security functions are working as expected; audits implementation and effectiveness of security controls.
Supervisor	Responsible for ensuring that users under their supervision comply with security requirements.
Users	Responsible for implementing security requirements at their level, which includes obeying policies and generally using good security hygiene.

accountable and responsible for the actions of the organization. Some roles, however, deal with the daily work of securing assets and implementing security controls. Table 1.3-1 describes some of these roles and related responsibilities.

Security Control Frameworks

Frameworks are overarching processes and methodologies that prescribe a path for the organization to perform security functions. There are risk management frameworks that recommend risk methodologies and steps to take to assess and respond to risk, just as there are frameworks used to manage an organization's IT assets. More closely related to the "in-the-weeds" security functions are security control frameworks.

TABLE 1.3-2 Commonly Used Security Control Frameworks

Framework	Description
National Institute of Standards and Technology (NIST) Special Publication 800-53	Security control framework promulgated by NIST; mandatory for U.S. federal government use and optional for all others. Consists of detailed security controls spanning areas such as access control, auditing, account management, configuration management, and so on.
International Organization for Standardization (ISO)/ International Electrotechnical Commission (IEC) 27002	Consists of information security controls used internationally and covers areas such as access control, physical and environmental security, cryptography, and operational security; part of the ISO/IEC 27000 series of standards covering information security management systems.
The Center for Internet Security (CIS) Controls	Consists of 18 controls (as of version 8, May 2021) in areas such as inventory and asset control, data protection, secure configuration, vulnerability management, and so on.
COBIT	Set of practices that are used to execute IT governance, including some security aspects. Note that the current version is COBIT 19.
Payment Card Industry (PCI) Data Security Standards (DSS)	Set of technical and operational controls established by the PCI Security Standards Council to protect cardholder data; consists of 15 Security Standards, as of version 3.2.1.

Security control frameworks prescribe formalized sets of controls, or security measures, an organization should implement to protect its assets and reduce risk. There are a variety of security control frameworks available for the organization to use, and many are mandated by external governance or internal adoption to align with the organization's mission and strategy. Table 1.3-2 lists and describes a few of the more commonly used frameworks.

Due Care/Due Diligence

The terms due care and due diligence are often used interchangeably but are not necessarily the same thing. However, you should know the difference for the exam, since sometimes the context of what you are discussing when using these two terms makes the distinction between them important. Their meanings are similar, and both are necessary to ensure that the organization is planning and doing the right things, for a variety of reasons. These reasons include avoiding legal liability, upholding reputation, and maintaining compliance with external governance, such as laws and regulations.

Due diligence means that the organization has put thought into planning its actions and has implemented controls to protect systems and information and prevent incidents such as a breach. Due diligence has more of a strategic focus. If the organization has carefully considered risk and implemented controls that protect systems, information, equipment, facilities, and people, then it is said to have practiced due diligence.

Due care is the term for actions the organization will take or has taken in response to a specific event. Due care typically applies to specific situations and scenarios, such as how an organization responds during a fire or natural disaster to protect lives and resources, or how it responds to a particular cybersecurity incident, such as an information breach. Because of this operational focus, it's often referred to as the "prudent person" rule, referring to what a reasonable person would do in the same or similar situation. Due care also involves verifying that the planning and actions taken as part of its due diligence responsibilities are practiced, effective, and work.



EXAM TIP Think of due diligence as careful planning and acting responsibly before something bad happens (proactive), and due care as acting responsibly when it *does* happen (reactive).

REVIEW

Objective 1.3: Evaluate and apply security governance principles In this objective we discussed security governance and its supporting concepts. We looked at both internal and external governance. Internal governance comes from the organization's own policies and procedures. External governance comes from laws and regulations. We also looked at how security functions integrate and align with the organization's strategy, goals, mission, and objectives. We discussed how organizational processes, such as acquisitions, divestitures, and so on, can both affect and are affected by security governance. We examined various organizational roles and responsibilities with regard to managing information technology and security. We also considered the need for security control frameworks and how they form the basis for protecting assets within the organization. Finally, we reviewed the key concepts of due care and due diligence and how they are necessary to reduce risk and liability for the organization.

1.3 QUESTIONS

1. The executive leadership in your company is concerned with ensuring that internal governance reflects its commitment to follow laws and statutes imposed on it by government agencies. Which of the following is used internally to translate legal requirements into mandatory actions organizational personnel must take in certain circumstances?
   - **A.** Standards
   - **B.** Strategy
   - **C.** Regulations
   - **D.** Policies

2. Which of the following does the information security strategy directly support?
   - **A.** Organizational mission
   - **B.** Organizational goals
   - **C.** Organizational business strategy
   - **D.** Operational plans
3. Which of the following senior roles has the responsibility for ensuring customer, organization, and employee data are kept secure and used properly?
   - **A.** Chief privacy officer
   - **B.** Chief security officer
   - **C.** Chief information officer
   - **D.** Data owner
4. Gail is a cybersecurity analyst who is contributing to the information security strategy document. The organization is going to expand internationally in the next five years, and Gail wants to ensure that the control framework used supports that organizational goal. Which of the following control frameworks should she include in the information security strategy for the organization to migrate to over the next few years?
   - **A.** NIST Special Publication 800-53
   - **B.** ISO/IEC 27002
   - **C.** COBIT
   - **D.** CIS Controls

1.3 ANSWERS

1. **D** Policies are used to translate legal requirements into actionable requirements that organizational personnel must meet.
2. **C** The organizational information security strategy directly supports the primary organizational business strategy, which in turn supports the goals of the organization and its overall mission.
3. **A** The chief privacy officer has responsibility for ensuring customer, organization, and employee data are kept secure and used properly. The chief security officer is responsible for all aspects of organizational security. The chief information officer is concerned with the entire IT infrastructure. A data owner is concerned with a particular type and sensitivity of data and is responsible for determining access controls for that data.
4. **B** Gail should include the International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 27002 control framework in the organization's information security strategy for implementation in the organization over the next several years, since it can be used internationally and is not tied to a particular government or business standard.



Objective 1.4 Determine compliance and other requirements

Directly following our discussion on governance, Objective 1.4 discusses the necessity for security programs to be compliant with that governance. In this objective we will look at the legal and regulatory aspects of obeying governance, as well as how governance also affects contractual agreements and privacy.

Compliance

In the previous objective we discussed governance. Think of this objective, regarding compliance, as a natural extension of that topic, since complying with governance requirements is a critical part of cybersecurity. *Compliance* means obeying the requirements of a particular governance standard. Remember that governance can be external or internal. External governance is usually in the form of laws, statutes, or regulations established by the government. The organization typically has no influence or control over the application of external governance. However, it does control its own internal governance. Internal governance comes in the form of the organization's own policies, procedures, standards, and guidelines.

Cross-Reference

Internal governance documents will be discussed further in Objective 1.7.

Compliance is monitored through a variety of means. An organization may be assessed by the agency issuing the requirements or through a third party that is authorized by the regulating agency. Typically, compliance is checked through the following:

- Inspections
- Audits
- Required reports
- Investigations

If an organization is deemed noncompliant with governance requirements, there may be short- or long-term consequences. In some cases, consequences could simply be an unfavorable report or a requirement to be compliant within a specified period of time. In other cases, the consequences of noncompliance can be quite severe. Potential consequences of noncompliance include

- Criminal charges and prosecution
- Legal liability

- Civil suits
- Fines
- Loss of stakeholder or consumer confidence

In this objective we will discuss compliance with several different types of requirements, including laws and regulations, contracts, and industry standards. We'll also talk about compliance with privacy rules, which are found in laws and other types of governance.

Legal and Regulatory Compliance

Most laws and regulations focus on requirements to protect specific categories of sensitive data. These categories include personally identifiable information (PII) (elements such as name, address, driver's license number, passport number, and so on), protected health information (PHI), and financial information. Not only do these laws specify requirements for administrative, technical, and physical security controls that must be implemented to protect data, but many of these laws are also designed to protect consumers and their privacy. Some laws even dictate how organizations must report and respond to data breaches. Table 1.4-1 offers a sampling of regulations that have cybersecurity ramifications or must be routinely considered by cybersecurity personnel for compliance.



EXAM TIP You do not have to know the particulars of any law or regulation for the exam, but you should be generally familiar with them for both the exam and your career.

TABLE 1.4-1 Laws and Regulations That Affect Cybersecurity

Law, Regulation, or Statute	Description
Federal Information Security Management Act (FISMA)	Directs all federal government agencies to manage risk and implement cybersecurity controls
Health Insurance Portability and Accountability Act (HIPAA)	Establishes requirements for protecting the security and privacy of protected health information (PHI)
Health Information Technology for Economic and Clinical Health (HI-TECH) Act	Expands the requirements of HIPAA to include penalties for noncompliance and requirements for breach notification
Gramm-Leach-Bliley Act of 1999	Imposes requirements on banks and other financial institutions to protect individual financial data
General Data Protection Regulation (GDPR)	European Union regulation implemented in May 2018 to protect the personal data and privacy of EU citizens
California Consumer Privacy Act (CCPA)	Far-reaching U.S. state law focused on protecting the PII of California residents, as well as breach notification
Sarbanes-Oxley (SOX) Act	Requires corporations to establish strong internal cybersecurity controls

Contractual Compliance

As data protection has increased in criticality over the past decade, many contracts between organizations are now including language describing how each organization will protect sensitive data shared between them. This language is more specific than a standard nondisclosure agreement (NDA). Contract terms now include responsibilities for data protection, data ownership, breach notification, and audit requirements. Companies entering these contractual relationships must be compliant with the terms or face civil liability for breach of contract.

Often contracts and other types of agreements require demonstration of due diligence and due care on the part of all entities entering into the contract. Sometimes contracts even outline penalties for violating the terms of the contract, such as failing to protect sensitive information. Sometimes this contract language is included as a regulatory requirement to protect specific categories of information, such as protected health information, that may be shared between two healthcare providers. Other times, a company may want to include contractual provisions to limit the liability of one party if the other party suffers a breach or incurs liability.

Compliance with Industry Standards

Some industries have their own professional standards bodies that establish governance regulations for that industry. If an organization desires to become a member of one of these professional standards bodies, it must adopt and agree to obey the regulatory requirements established by that standards body for the industry. Noncompliance with those standards may result in the organization no longer being sanctioned by the professional organization, or even forbidden to further participate in activities governed by that body. An example of an industry standard is the Payment Card Industry (PCI) Data Security Standard (DSS), which is a security standard developed and enforced by a consortium of credit card providers, such as Visa, MasterCard, American Express, and others. PCI DSS is mandated for any merchant that processes credit card payments. As the standard is voluntary, a merchant does not have to agree to comply, but the consequence is that the merchant will not be permitted to process credit card transactions.

Privacy Requirements

There are many different laws, regulations, and even industry standards that cover privacy. Remember that privacy is different from security in that security seeks to protect the confidentiality, integrity, and availability of information, whereas privacy governs what is done with specific types of information, such as PII, PHI, personal financial information, and so on. Privacy determines how much control an individual has over their information and what others can do with it, to include accessing it and sharing it. You can think of privacy as controlling how information is used and security as the mechanism for enforcing that control.

We briefly mentioned a few of the most prevalent privacy regulations earlier in the objective. Although there are differences in how countries view privacy and enforce privacy rules, the privacy laws and other governance standards of most countries have some common elements.

Whether it is the General Data Protection Regulation enforced by the European Union or the NIST Special Publication 800-53 privacy controls, there are some commonalities in the different privacy requirements. Complying with privacy laws and regulations usually requires an organization that collects an individual's (the *subject's*) personal data to have a formal written privacy policy, and then further demonstrate how it complies with that policy. Privacy policies typically include, at a minimum, the following provisions:

- **Purpose** The purpose for which information is collected
- **Authority** The authority to collect such information
- **Uses** How the collected information is used
- **Consent** The extent of the subject's consent required to collect and use information
- **Opt-in/opt-out** Rights of the subject to opt out of data collection
- **Data retention and disposal** How long the information is retained and how it is disposed of
- **Access and corrections** How subjects of the information can view and correct their information
- **Protection of privacy data** How the organization protects the information and who is responsible for that protection
- **Transfer to third parties** The circumstances under which the information can be released to third parties
- **Right to be forgotten** The guarantee that an individual can request to have their information deleted
- **Notification of breach** The requirement to notify the subject of any breach of their personal informaion

Note that some of these privacy requirements are not always applicable, depending upon the law, regulation, or even country involved. In Objective 1.5 we will discuss additional privacy requirements and concerns, including specifics on how privacy is treated on an international basis.

REVIEW

Objective 1.4: Determine compliance and other requirements This objective covered the necessity of complying with governance requirements. Compliance with laws, regulations, contracts, industry standards, and privacy requirements is a major portion of cybersecurity. First, we discussed compliance with several different laws and regulations imposed by governments. Laws and regulations primarily serve to enforce how particular categories of data are protected, such as financial, healthcare, and personal data. Compliance with laws and regulations is mandatory, and lack of compliance is typically punished by fines and civil penalties, but some laws have provisions that specify possible criminal penalties such as imprisonment.

We also discussed another aspect of civil penalty—contract compliance. Contracts are agreements between two or more entities and are legally enforceable. Failure to comply with the terms of a contract can result in civil liabilities, such as lawsuits and fines.

Although industry standards may not be legally mandated, participation in a particular industry may require that an organization obey those standards. A classic example is the security standards imposed by the credit card industry, known as the Payment Card Industry Data Security Standard (PCI DSS), which dictates how organizations that process credit card payments must secure their systems and data.

Finally, we examined common characteristics of privacy rule requirements in several laws, regulations, and other governance standards. These rules include an individual's ability to be able to correct erroneous information, determine who has access to personal information, and the right to be informed of a breach of personal data.

1.4 QUESTIONS

1. Emma is concerned that the recent breach of personal health information in a large healthcare corporation may affect her, but she has not yet been notified by the company that was breached. Emma, a resident of the state of Alabama, is researching the various laws under which she should be legally notified of the breach. Which of the following relevant laws or regulations dictates the timeframe under which she should be notified of the data breach of her PHI?
   - **A.** California Consumer Privacy Act (CCPA)
   - **B.** Health Information Technology for Economic and Clinical Health (HI-TECH) Act
   - **C.** General Data Protection Regulation (GDPR)
   - **D.** Federal Information Security Management Act (FISMA)
2. Riley is a junior cybersecurity analyst who recently went to work at a major banking institution. One of the senior cybersecurity engineers told him that he must become familiar with the different data protection regulations that apply to the financial industry. With which of the following laws or regulations must Riley become familiar?
   - **A.** General Data Protection Regulation (GDPR)
   - **B.** Federal Information Security Management Act (FISMA)
   - **C.** Gramm-Leach-Bliley Act of 1999
   - **D.** Health Insurance Portability and Accountability Act (HIPAA)
3. Geraldo owns a small chain of sports equipment supply stores. Recently, his business was required to undergo an audit to measure compliance with the PCI DSS standards. Geraldo's business failed the audit. Which one of the following is the most likely consequence of this failure?
   - **A.** His business may no longer be allowed to process credit card transactions unless he remediates any outstanding security issues.
   - **B.** His business will be required to report ongoing compliance status under FISMA.

C. His business must report a data breach under HIPAA.

D. His business will face potential fines under the HI-TECH Act.

4. Nichole is a contracts compliance auditor in her company. She is reviewing the cybersecurity requirements language that should be included in a contract with a new business partner. The new partner will have access to extremely sensitive information owned by Nichole's company. Which of the following is critical to include in the contract language?

A. The requirement for the business partner to maintain high-availability systems

B. The requirement for the business partner to immediately notify Nichole's company in the event the partner suffers a breach

C. The business partner's legal obligations under the law

D. The business partner's security plan

1.4 ANSWERS

1. **B** Emma should be notified of the breach under the Health Information Technology for Economic and Clinical Health (HI-TECH) Act, which expands HIPAA regulations to include breach notification. As a resident of the state of Alabama, neither the California Consumer Privacy Act (CCPA), which protects state of California residents, nor the General Data Protection Regulation (GDPR), which protects citizens of the European Union, applies. FISMA is a federal regulation requiring government agencies to manage risk and implement security controls.

2. **C** Riley must become familiar with and understand the requirements imposed by the Gramm-Leach-Bliley Act of 1999, which requires financial institutions to implement proper security controls to store, process, and transmit customer financial information.

3. **A** Because Geraldo's business failed an audit under the Payment Card Industry Data Security Standard, his business could potentially be banned from processing credit card transactions until the issues are remediated.

4. **B** Since the business partner will have access to extremely sensitive information, Nichole should include language in the contract that requires the partner to immediately notify her company if there is a data breach. High-availability requirements for the business partner are not relevant to protecting sensitive data. Nichole does not have to include the business partner's obligations under the law in the contract language, since the law applies whether or not the language is in the contract. The security plan would not normally be included in contract language.



Objective 1.5 Understand legal and regulatory issues that pertain to information security in a holistic context

In Objective 1.5 we are going to continue our discussion regarding legal and regulatory requirements an organization may be under for governance, compliance, and information security in general. We will examine issues such as legal and regulatory requirements, cybercrime, intellectual property, and transborder data flow, as well as import/export controls and privacy issues.

Legal and Regulatory Requirements

As a cybersecurity professional, you should understand the legal and regulatory requirements imposed on the profession with regard to protecting data, prosecuting cybercrime, and dealing with privacy issues. You also should be familiar with the issues involved with exporting certain technologies to other countries and how different countries view data that crosses or is stored within their borders. Many of these issues are interrelated, as countries enact data laws that benefit their own national interests while simultaneously affecting privacy, technologies that can be used within their borders, and how data enters and exits their nation's IT infrastructures.

Cybercrimes

The definition of what constitutes a cybercrime varies by country, but in general, a *cybercrime* is a violation of a law, statute, or regulation that is perpetrated using or targeting computers, networks, or other related technologies. Common cybercrimes include hacking, identity theft, fraud, cyberstalking, child exploitation, and the propagation of malicious software. The CISSP exam does not expect you to be an expert on law enforcement, but you should be familiar with some of the current laws and issues related to cybercrime. These include data breaches and the theft or misuse of intellectual property.

Cross-Reference

Areas related to cybercrime and cyberlaw, such as investigations, are covered in Objectives 1.6 and 7.1.

Data Breaches

Data theft, loss, destruction, and access by unauthorized entities has now become the largest concern in the cybersecurity world. Data breaches are now commonplace, because the value of sensitive data has motivated sophisticated individuals and gangs to expend a lot of time and resources toward attacking computer systems. Adding that to the fact that often inadequate protections

may be sometimes put in place to protecting sensitive data. Although slow to catch up with the fast-moving pace of cybercrime, data breach laws have been put in place to deter such instances, as well as to deter such instances by imposing heavy penalties and by giving the legal system more leeway to investigate, prosecute, and punish those who carry out these crimes.

Some data breach laws apply to specific areas, such as healthcare information, financial data, or personal information. Others apply across the board regardless of data type. Typically, data breach laws define the types of data they are attempting to protect and specify penalties to be imposed on the perpetrator of a breach. Data breach laws include breach notification provisions that require an organization that suffers a data breach to notify subjects potentially impacted by the breach, usually within a specified time period, as well as impose fines and penalties for inadequate data protection or failure to notify subjects in case of a breach. Various U.S. laws that address data protection requirements, as well as data breach concerns, include

- Health Information Technology for Economic and Clinical Health (HI-TECH) Act
- California Consumer Privacy Act (CCPA)
- Economic Espionage Act of 1996
- Gramm-Leach-Bliley Act of 1999

Licensing and Intellectual Property Requirements

Intellectual property (IP) refers to tangible and intangible creations by individuals or organizations, such as expressions, ideas, inventions, and so on. Intellectual property can be legally protected from use by anyone not authorized to do so. Examples of intellectual property include software code, music, movies, proprietary information, formulas, and so forth. The different types of intellectual property you should be familiar with for the exam include

- **Trade secret** Confidential information that is proprietary to a company, which gives it a competitive advantage in a market space. Trade secrets are legally protected as the property of their owners.
- **Copyright** Protections for the rights of creators to control the public distribution, reproduction, display, adaptation, and use of their own original works. These works could include music, video, pictures, books, and even software. Copyright applies to the *expression* of an idea, not the idea itself. Note that copyright exists at the moment of expression—registration provides legal protection but is not required for ownership, although highly recommended for ownership enforcement.
- **Trademark** A word, name, symbol, sound, shape, color, or some combination of these things that represents a brand or organization. Trademarks are typically distinguishing marks that are used to identify a company or product.
- **Patent** The legal registration of an invention that provides its creator (or owner if the owner is not the same entity as its creator) with certain protections. These protections include the right of the invention owner/creator to determine who can legally use the invention and under which circumstances.

While trade secrets are not typically registered (due to their confidential nature), their owners can initiate legal action against another entity that uses their proprietary or confidential information if the owner can prove their ownership. The other three forms of intellectual property are typically registered with a regulatory body to protect the owner's legal rights. Copyrights are not required to be registered, if the owner can prove that they are the owner/creator of a work if it is copied or used without their permission. A person desiring to legally use material protected by copyright, trademark, and patent laws must obtain a license from the owner of the intellectual property. The license describes the conditions under which an entity may legally use someone else's IP and gives them legal permission to do so, often for a fee.



EXAM TIP Of the intellectual property types we have discussed; trade secrets are not normally registered with anyone, unlike copyrights, trademarks, and patents, due to their confidential nature. However, if someone violates another organization's trade secrets, the entity claiming ownership to the trade secret should be able to prove that it belongs to them.

Import/Export Controls

Many countries restrict the import or export of certain advanced technologies; in fact, some countries consider importing or exporting some of these advanced technologies to be equivalent to importing or exporting weapons. Import/export controls that cybersecurity professionals need to be aware of specifically include those related to encryption technologies and advanced high-powered computers and devices. Each country has its own laws and regulations governing the import and export of advanced information technologies. The following are two key United States laws that address the export of prohibited technologies:

- **International Traffic in Arms Regulations (ITAR)** Prohibits export of items designated as military and defense items
- **Export Administration Regulation (EAR)** Prohibits the export of commercial items that may have military applications

Consider the impact if advanced encryption technologies were to fall into the hands of a terrorist or criminal organization, or the declared enemy of a country. Obviously, countries operate in their own best interests when declaring which technologies may or may not be imported or exported to or from them. Another example would be a country that does not permit advanced encryption technologies to be imported and used by its citizens, because the government wants to outlaw encryption methodologies that it cannot decrypt.

The Wassenaar Arrangement is an international treaty, currently observed by 42 countries, that details export controls for specific categories of dual-use goods. Of interest to cybersecurity personnel are the Category 3 (Electronics), 4 (Computers), and 5 (Telecommunications and Information Security items) areas, which should be consulted prior to export, based upon the laws of both the exporting and importing countries.

Transborder Data Flow

Data flow between the borders of countries is sometimes subject to controversy. Issues to consider include privacy concerns, proprietary data, and, of course, data that could be classified as confidential by a foreign government. Some countries strictly prohibit certain classifications of data from being processed or even stored in a different country. These laws are called *data localization* laws (also called *data sovereignty* laws) and require that certain types of data be stored or processed within a country's borders.

In some geographical areas, such as the European Union, the prohibitions on processing the private data of its citizens are intended to protect people. In other cases, a country may require data to stay within its borders because of the desire to restrict, control, and access such data, due to censorship, national security, or other motivation. While there are different international agreements that control some of these cross-border data flows, there are also social, political, and technological concerns over how effective these agreements are when placed in the context of different privacy rights and technologies such as cloud computing, as well as the globalization of IT.

Privacy Issues

Privacy can be a complicated issue, particularly when discussing it in the context of international laws and regulations. In some areas of the world, such as the European Union, privacy is a priority and is strictly enforced. In other locales, the adherence to any semblance of personal privacy is essentially lip service. Countries often define their privacy laws in relation to several other issues, such as national security, data sovereignty, and transborder data flow. Some countries have specific laws and regulations that are enacted to protect personal privacy, such as:

- European Union's General Data Protection Regulation (GDPR)
- Canada's Personal Information Protection and Electronic Documents Act
- New Zealand's Privacy Act of 1993
- Brazil's Lei Geral de Proteção de Dados (LGPD)
- Thailand's Personal Data Protection Act (PDPA)

Other countries, including the United States, have no specific overarching privacy law, but tend to include privacy requirements as part of other laws, such as those that affect businesses or a specific market segment or population. Examples of this approach in the United States include the Gramm-Leach-Bliley Act (GLBA) of 1999, which applies to financial organizations, the Health Insurance Portability and Accountability Act (HIPAA), levied on healthcare providers, and the Privacy Act of 1974, which applies to only U.S. government organizations processing privacy information.

In addition to the different privacy policy elements we discussed in Objective 1.4, such as purpose, authority, use, and consent, there are different methods of addressing privacy in law and regulation. One way is within a particular industry or market segment, called *vertical enactments.* Privacy laws and regulations are enacted and apply to a specific area, such as the

healthcare field or the financial world (e.g., HIPAA and GLBA, respectively). Contrast this to a *horizontal enactment*, where a particular law or regulation spans multiple industries or areas, such as those laws that protect PII, regardless of its industry use or context.

REVIEW

Objective 1.5: Understand legal and regulatory issues that pertain to information security in a holistic context In Objective 1.5 we continued the discussion of compliance with laws and regulations and delved into critical cybersecurity issues, such as cybercrime, data breaches, theft of intellectual property, import and export of restricted technologies, data flows between countries, and privacy. Cybercrime is a violation of a law, statute, or regulation that is perpetrated using or targeting computers, networks, or other related technologies. Common cybercrimes include hacking, identity theft, fraud, cyberstalking, child exploitation, and the propagation of malicious software. A data breach is theft or destruction of data, typically through a criminal act. Several laws have been enacted to deal with breaches of specific kinds of data, including those applicable to both the healthcare and financial industries.

We also discussed the different types of intellectual property that must be protected, including trade secrets, copyrights, trademarks, and patents. Trade secrets are legally protected but are not typically registered due to their confidential nature. Copyrights also do not have to be registered but should be to protect their owners' legal interests. Trademarks and patents are legally registered with an appropriate government agency. A license is required for someone to legally use someone else's IP protected by copyright, trademark, or patent laws.

Import and export controls are designed to prevent certain advanced technologies, such as encryption, from entering or leaving a country's borders, based on the country's own laws and regulations. Several treaties have been enacted between countries restricting the import or export of certain sensitive technologies, including the Wassenaar Arrangement. Transborder data flow is subject to the laws and restrictions of different countries, based on their own national self-interests. Data localization or sovereignty laws are imposed to restrict the export, use, and access to certain categories of sensitive data. Privacy issues are compounded by the lack of consistency in international laws and the lack of respect for individual privacy in certain countries.

1.5 QUESTIONS

1. Which of the following laws requires breach notification of protected health information (PHI)?
   - **A.** HI-TECH
   - **B.** GLBA
   - **C.** PCI DSS
   - **D.** CCPA

2. In order for crime to be considered a cybercrime, which of the following must be true?
   - **A.** It must result in fraud.
   - **B.** It must use computers, networks, and/or related technologies.
   - **C.** It must involve malicious intent.
   - **D.** It must not be a violent crime.
3. Your company has produced a secret formula used to manufacture a particularly strong metal alloy. Which of the following types of intellectual property would the secret formula be considered?
   - **A.** Trade secret
   - **B.** Trademark
   - **C.** Patent
   - **D.** Copyright
4. A country bans importation of high-strength encryption algorithms for use within its borders, since it desires to be able to intercept and decrypt messages sent and received by its citizens. Which of the following laws might it enact to restrict these technologies from being used?
   - **A.** Copyright laws
   - **B.** Intellectual property laws
   - **C.** Privacy laws
   - **D.** Import/export laws

1.5 ANSWERS

1. **A** The Health Information Technology for Economic and Clinical Health (HI-TECH) Act is a law enacted to further protect private healthcare information and provides for notification to the subjects of such information if it has been breached.
2. **B** A crime is considered a cybercrime if it targets computers, networks, or related technologies, regardless of the intent, whether fraud is committed, or whether the crime results in physical violence.
3. **A** Because the formula is considered confidential and gives the company an edge in the market, it would be considered a trade secret. The formula would not be registered under copyright, trademark, or patent laws, because this would divulge its contents to the public.
4. **D** If a country wishes to restrict the use of advanced technologies, such as encryption, by its citizens and within its borders, it will enact import/export laws to prevent those technologies from entering the country and make their use or possession illegal.



Objective 1.6

Understand requirements for investigation types (i.e., administrative, criminal, civil, regulatory, industry standards)

In Objective 1.6 we will discuss investigations, and examine the various investigation types, such as administrative investigations, as well as civil, criminal, and regulatory ones. We will also look at various industry standards for investigations that may not fall into one of the other categories.

Investigations

Investigations are a necessary part of the cybersecurity field. Frequently, investigations are conducted because someone doesn't obey the rules, such as those found in acceptable use policies, or someone makes a mistake that results in data compromise or loss. Regardless of the reason that prompts the investigation, a cybersecurity professional should be familiar with the different types of investigations that may be needed. Note that this objective discusses the different investigation types; it is a valuable prerequisite for the much more detailed discussion of investigations that we will have later in a related objective in Domain 7.

Cross-Reference

Investigations are also covered in Objective 7.1.

Administrative Investigations

An administrative investigation is one that focuses on members of an organization. This type of investigation usually is an internal investigation that examines either operational issues or a violation of the organization's policies. Administrative investigations are usually conducted by the organization's internal personnel, such as cybersecurity personnel or even auditors. In small organizations, management may designate someone to conduct an independent investigation or even consult with an external agency. Consequences resulting from an internal administrative investigation include, for example, reprimands and employment termination. Sometimes, however, the investigation can escalate into either a civil or criminal investigation, depending on the severity of the violations.

Civil Investigations

A civil investigation typically occurs when two parties have a dispute and one party decides to settle that disagreement by suing the other party in court. As part of that lawsuit, an investigation is often necessary to establish the facts and determine fault or liability. Based on which

party the court deems liable, the party at fault may incur fines or owe money (damages) to the party considered harmed. Note that the evidentiary requirements (burden of proof) of civil investigations are not as stringent as the evidentiary requirements of criminal investigations. Civil investigations use a "preponderance of the evidence" standard, meaning that the case could be decided based on just a reasonable possibility that someone committed a wrongdoing against another party. Note that regardless of the burden of proof requirements levied on a civil versus a criminal investigation, it does not change the conduct of the investigation, as we will see later on in Objective 7.1.

Criminal Investigations

More serious investigations often involve circumstances where an individual or organization has broken the law. Criminal investigations are conducted for alleged violations of criminal law. Unlike administrative investigations, criminal investigations are typically conducted by law enforcement personnel. As previously noted, the standard of evidence for criminal investigations is much higher than the standard for civil investigations and requires a determination of guilt or innocence "beyond a reasonable doubt," since the penalties are much more serious. Penalties that could result from a criminal investigation and subsequent trial include fines or imprisonment.



EXAM TIP The primary differences between civil and criminal investigations are that civil investigations are part of a lawsuit, and the burden of proof is much lower than in a criminal investigation. Civil cases also usually have less severe penalties than criminal cases.

Regulatory Investigations

A regulatory investigation may be conducted by a government agency when it believes an individual or organization has violated administrative law, typically a regulation or statute meant to control the behavior of organizations with regard to societal responsibility, due care and diligence, or economic harm toward others. Unlike a criminal investigation, a regulatory investigation, however, does not necessarily have to be conducted by law enforcement personnel. It can be conducted by other government agencies responsible for enforcing administrative laws and regulations.

An example of a regulatory investigation is one where the Securities and Exchange Commission (SEC) investigates a company for insider trading. The penalties imposed by regulatory investigations can range from the same penalties received after a civil investigation, such as fines or damages, or even to those resulting from a criminal investigation, such as imprisonment.



EXAM TIP The primary difference between a criminal investigation and a regulatory investigation is context. Criminal investigations may occur if there is obvious serious evidence of fraud, violence, or another serious crime. Criminal investigations also more often than not involve individuals. Regulatory investigations are normally conducted when an organization (versus an individual) breaks an administrative type of law, such as laws related to financial crimes or data protection. A regulatory investigation can easily turn into a criminal one, depending upon the circumstances.

Industry Standards for Investigations

The final type of investigation we will discuss is not imposed by laws or regulations or even by any specific organization. An industry standard is imposed by organizations within the industry itself, as a way for an industry to self-regulate the behavior of its members, whether they are individuals or organizations. Industry standards are usually voluntary, and an organization typically adopts them or must agree to them to be part of that industry.

An example of an industry standard is the Payment Card Industry Data Security Standard (discussed in the previous two objectives), which all vendors and merchants that process credit card transactions must obey or risk losing their payment card processing privileges. Another industry standard that applies more directly to individual cybersecurity personnel is the standard imposed by (ISC)^2 to be certified as a CISSP. An individual agrees to the professional and ethical standards as a condition of certification.

Organizations that agree to abide by industry standards also agree to be investigated in the event they are suspected of violating the standards. An investigation by a standards organization normally is conducted for cause, such as a complaint filed against a member organization, rather than as part of a routine audit process. The investigations are usually carried out by members of the standards body, and the penalties that may be imposed on organizations and individuals found violating industry standards may range from suspension, termination from the standards body, fines, or censure. In extreme cases, a permanent ban from the organization sponsoring the standard or even civil liabilities may occur.

REVIEW

Objective 1.6: Understand requirements for investigation types (i.e., administrative, criminal, civil, regulatory, industry standards) In this objective you learned about the different types of investigations, including administrative, civil, criminal, regulatory, and those required by industry standards. Administrative investigations are conducted within an organization by internal security or audit personnel. Civil investigations are conducted as part of a lawsuit between parties and are designed to determine which party is at fault. Criminal investigations are conducted when a person or organization has broken the law and may result in stiff penalties imposed on the guilty party, such as fines

or imprisonment. Regulatory investigations are conducted by agencies responsible for enforcing administrative laws and statutes. Regulatory investigations can also result in fines, damages, or imprisonment. Violating a standard imposed by an industry or professional organization can result in an investigation by an enforcing body. Penalties resulting from this type of investigation could include suspension from the organization, termination from the industry, or fines.

1.6 QUESTIONS

1. You are a cybersecurity analyst in a medium-sized company and have been tasked by your senior managers to investigate the actions of an individual who violated the organization's acceptable use policy by accessing prohibited websites. During the investigation, you determine that the individual's Internet access also potentially violated laws in the state where the company is located. Your management makes the decision to turn the investigation over to law enforcement authorities. Which of the following best describes this type of investigation?
   - **A.** Administrative investigation
   - **B.** Civil investigation
   - **C.** Regulatory investigation
   - **D.** Criminal investigation
2. One of your company's web servers was hacked recently. After your company investigated the hack and mitigated the damage, another company claimed that the attacker used your company's web server to attack its network. The other company has initiated a lawsuit against your company and has hired a private cybersecurity investigation firm to determine if your company is liable. Which of the following types of investigation would this be?
   - **A.** Criminal investigation
   - **B.** Administrative investigation
   - **C.** Civil investigation
   - **D.** Investigation resulting from violating an industry standard
3. Your company has joined an industry professional organization, which imposes requirements on its member organizations as a condition of membership. A competitor recently reported your company to the professional organization for violating its rules of behavior. The professional organization has decided to launch an independent investigation to validate these claims. What type of investigation would this be considered?
   - **A.** Administrative investigation
   - **B.** Civil investigation
   - **C.** Industry standards investigation
   - **D.** Criminal investigation

4. Which of the following examples best describes a regulatory investigation?
   - **A.** A company's cybersecurity team investigates violations of acceptable use policy.
   - **B.** Corporate lawyers investigate allegations of trademark infringement by another corporation.
   - **C.** The Federal Bureau of Investigation investigates allegations of terrorist support activities by individuals in your company.
   - **D.** The Federal Communications Commission investigates allegations of unlawful Internet censorship by an Internet service provider.

1.6 ANSWERS

1. **D** Although the investigation started as a simple internal administrative investigation, the discovery of a potential violation of the law escalated the investigation into a criminal investigation since law enforcement authorities have been called in.
2. **C** Since the investigation was initiated as the result of a civil lawsuit, this would be considered a civil investigation.
3. **C** Since the investigation was initiated because of a claim that your company has violated the requirements imposed by an industry standards organization, this would be considered an industry standards investigation.
4. **D** The Federal Communications Commission (FCC) investigating potential unlawful censorship by an Internet service provider would be an example of a regulatory agency investigation. A crime investigated by the FBI would be considered a criminal investigation. Corporate lawyers investigating trademark infringement would constitute a civil investigation. Cybersecurity personnel within an organization investigating violation of acceptable use policy would be considered an administrative investigation.



Objective 1.7 Develop, document, and implement security policy, standards, procedures, and guidelines

Objective 1.7 will close out our discussion on governance. For this objective we will look at internal governance, such as security policy, standards, procedures, and guidelines. These are internal governance documents developed to support external governance, such as laws and regulations.

Internal Governance

As briefly discussed in Objective 1.3, internal governance exists to support and articulate external governance that may come in the form of laws, regulations, statutes, and even professional industry standards. Internal governance specifically requires internal organizational personnel to support external governance, as well as the strategic goals and mission of the organization. Internal governance comes in the form of policies, procedures, standards, guidelines, and baselines, which we will discuss throughout this objective.

Internal governance is formally developed and approved by executive management within the organization. However, in all practicality, many line workers and middle managers also have input into internal governance. Often, they help provide the information or draft documents that senior managers finalize and approve. Internal governance is also often managed by an internal executive or steering committee, which is represented by a broad range of business areas within the organization, including business processes, IT, cybersecurity, human resources, and financial departments. This broad approach allows all important stakeholders to have a voice in the internal governance structure. Ultimately, however, senior management is responsible for approving and implementing all internal governance.

Policy

Security policy represents the requirements the senior leadership of an organization imposes on its security management program, and how it conducts that program. Individual security policies make up that overarching policy, and are the cornerstone of internal governance. Policies provide direction to organizational personnel. Policies dictate requirements that organizational personnel must meet. Note that policies and other internal governance documents are considered administrative controls. Policies don't go into detail—they simply state requirements. Most policies also list the roles and responsibilities of those who are required to manage or implement the policies. A policy dictates a requirement; it states *what* must done, and sometimes it even states *why* it must be done (to implement a law or regulation, for instance). However, it usually does not dictate *how* the requirement must be carried out. The process of *how* is described in the procedures, which we will discuss in the next section.

While organizations write their policies in many ways, policies should generally be brief and concise. A policy document should state a specific scope and purpose for the policy, and when tied to external governance, a policy should state which law or regulation it supports. A policy document may also state the consequences of not obeying the policy. Finally, a senior executive should sign the policy document as an approval authority for the policy, as this demonstrates management's commitment to the policy.

Procedures

Procedures, as well as other internal governance documents, exist to support policies. Where a policy dictates *what* must be done, a procedure goes into further detail and describes *how*

it must be done. A procedure can be quite detailed and describe the different processes and activities that must be performed to carry out the requirements of the policy. Procedures are often developed at a lower level in the organization, usually with middle management and line workers involved in their creation. Ultimately, they still must be approved by senior managers, but those managers are typically less involved in the actual writing of the procedure.

Procedures can detail a wide variety of processes, such as handling equipment, encrypting sensitive data, performing data backups, disposing of media, and so on. Note that, like policies, procedures are usually mandatory requirements in the organization. Procedures are often informed by standards and guidelines documents, discussed in turn next.

Standards

Standards can come in many forms. A standard may be a control framework, for example, or a document that describes the level of performance a procedure or process must attain to be considered performed properly. It also may detail minimum requirements for a process or activity. A standards document is usually a mandatory part of internal governance, just as policies and procedures are. A standard may be produced by an independent organization or a government regulatory agency. In some cases, an organization does not have a choice when it comes to adopting a standard. In other cases, the organization may choose to adopt a voluntary standard but make it mandatory for use across the organization. In any event, standards exist to provide direction on how procedures are performed.

To give you an idea of how standards relate within the internal governance framework, a policy may be created that mandates the use of security controls. It also, due to external governance, may mandate that the NIST security control catalog (NIST Special Publication 800-53, Revision 5, a standard mandatory for U.S. government entities) be used in all processes and procedures. Procedures may detail how to implement specific controls mandated in the NIST control catalog. Another example is when a policy mandates the use of encryption for data stored on sensitive devices. A procedure will detail the steps a user must take to encrypt data, and the Federal Information Processing Standards (FIPS) may dictate the requirements for the encryption algorithms used.

Guidelines

Guidelines are typically supplemental to standards and procedures. Guidelines can be developed internally by the organization, or they may be developed by a vendor or professional security organization. Guidelines are usually not considered mandatory since they only provide supplemental information on how to perform procedures or activities. A guideline could explain how to perform a task in greater detail or just provide additional information that may not be included in procedures. Sometimes guidelines provide best practices that are not considered mandatory but may be necessary.



EXAM TIP To help understand the relationships between policy, procedures, standards, guidelines, remember that guidelines are not mandatory, but the other elements of internal governance are. Also remember that policy is directive, procedures detail how to implement policy, standards dictate to what level or depth, and guidelines are simply supplemental information that can be of assistance in implementing policies.

Baselines

Like the previous internal governance documents we reviewed, a baseline is developed to implement requirements established by policy. Unlike those other documents, though, baselines are implemented as configuration items on different components within the infrastructure. A baseline is a standardized configuration used across devices in the organization. It could be a standardized operating system installation configured identically with other systems, or it could be standard applications consistently configured in a like manner. Baselines could also consist of standardized network traffic patterns.

The key factor about baselines is that they are standardized across the organization. They support security policies by translating policy into actual control implementation. For example, if a policy states that encryption for data at rest will be used across all infrastructure devices, and a standard states that it must be AES 256-bit encryption, the baseline will include configuration options to implement that requirement. The procedures would provide the details on how to configure that baseline.

Baselines are maintained and recorded as part of configuration management. When the organizational infrastructure changes, requiring a baseline change, this modification must be carefully planned, tested, executed, and recorded. Any accepted changes become part of the new baseline as part of formal change and configuration management procedures. Ultimately, all changes must be in compliance with the approved and implemented internal governance.



EXAM TIP Although baselines are not included as part of Objective 1.7, they are critical in understanding how organizational policy is implemented at the system and infrastructure level.

REVIEW

Objective 1.7: Develop, document, and implement security policy, standards, procedures, and guidelines In this objective we covered many components of internal governance. Internal governance reflects senior management leadership philosophy, as well as alignment with and support of external governance, such as laws and regulations.

Internal governance components include individual security policies, which state the requirements imposed by senior management on the organization to support its overarching security policy. Procedures detail how policies will be implemented, in terms of processes or activities. Standards help inform the degree of depth, quality, or level of performance that activities and processes must meet. Policies, procedures, and standards are typically considered mandatory. Guidelines are usually considered optional and consist of supplemental information used to enhance procedures with best practices or optimized methods of implementation. Guidelines can be developed by software or hardware vendors, professional organizations, or even the organization itself.

Baselines are the result of policy implementation and consist of standardized configurations for the infrastructure. Baselines include standardized operating systems, applications, network traffic, and security configurations. If the organization requires changes to the infrastructure due to new technologies, changes in business processes, or the threat landscape, these changes are incorporated into the baseline through change and configuration management processes.

1.7 QUESTIONS

1. Which of the following is a component of internal governance?
   - **A.** Laws
   - **B.** Regulations
   - **C.** Statutes
   - **D.** Policies
2. Which of the following is *not* considered a mandatory component of internal governance?
   - **A.** Guidelines
   - **B.** Standards
   - **C.** Policies
   - **D.** Procedures
3. You are a cybersecurity analyst in a medium-sized company. The senior management in your company, after a risk assessment, has decided to implement a policy that requires critical patches be applied to systems within one week of their release. Which of the following would detail the activities needed to implement that policy?
   - **A.** Operating system guidelines
   - **B.** Patch management procedures
   - **C.** A configuration management standard
   - **D.** A NIST-compliant operating system baseline

4. Your company has standardized baselines across the infrastructure for operating systems, applications, and network ports, protocols, and services. Recently, a new line-of-business application was installed but is not functioning properly. After examining the infrastructure security devices, you discover that one of the application's protocols and its associated port is blocked. What must be done, from a management perspective, to enable the application to work properly?
   - **A.** Uninstall the new line-of-business application, since its port and protocol are not allowed in the baseline
   - **B.** Go through the change and configuration management process to make the changes to the network traffic port to create a new permanent baseline
   - **C.** Unblock the associated protocol and port in the security device
   - **D.** Reconfigure the application so that it uses only ports and protocols already included in the baseline

1.7 ANSWERS

1. **D** Policies are used to implement internal governance requirements, and may align with external governance, such as laws, regulations, and statutes.
2. **A** Guidelines consist of supplemental information and are not considered mandatory parts of internal governance. They serve to enhance internal governance by providing additional information and best practices. Policies, procedures, and standards are considered mandatory components of internal governance.
3. **B** Patch management procedures would need to be updated after the policy change to include the requirement to implement critical patches to all systems within one week of their release. The procedures would detail exactly how these tasks and activities would be carried out.
4. **B** You should go through the formal change and configuration management process to add the application's port and protocol to the established baseline. Uninstalling the application is likely not an option, since the business decision was made to install it based on a valid business need. Simply unblocking the port and protocol the application uses on the security device is a technical approach, and may happen after the change to the baseline has been formally approved, but is not a management action. Reconfiguring the application may not be an option, since it likely uses specific ports and protocols for a reason and changing it may interfere with other applications on the network as well as create too many other changes to the baseline.



Objective 1.8 Identify, analyze, and prioritize Business Continuity (BC) requirements

Objective 1.8 begins a discussion that we will have throughout the book, through various other objectives, on business continuity planning (BCP). We will also discuss business continuity in Domain 7, and its closely related process, disaster recovery. For now, we will look at business continuity requirements such as those that are developed when performing a business impact analysis (BIA).

Business Continuity

Business continuity (BC) is a critical cybersecurity process that directly addresses the availability goal of security. BC is concerned with keeping the critical business processes up and running, even through major incidents, such as disasters and catastrophes. Although often discussed as a separate entity entirely, BC is intricately connected to incident response; BC is the process that often comes after the immediate concerns of containing and mitigating a serious incident and deals with bringing everything back up to its full operational status. BC is also closely related to disaster recovery; sometimes there is a blurry line where incident response, business continuity, and disaster recovery begin and end. While business continuity is concerned with keeping the business up and running, disaster recovery, as we will see later in Domain 7, focuses on the immediate concerns of safety, preserving human life, and recovering the equipment and facilities, so that business continuity can begin.



EXAM TIP Incident response, business continuity, and disaster recovery are three closely related but separate processes. Incident response is what immediately happens during any kind of a negative event, to discover what happened, how it happened, and how to stop the compromise of information and systems. Disaster recovery may also occur during that process, depending upon the nature of the incident, or it may be a separate process, but it is chiefly concerned with saving lives and equipment. Business continuity immediately follows disaster recovery and focuses on getting the business back into operation performing its primary mission. All of these activities, however, require integrated planning in advance of an event.

BC is also an integral part of risk management, as you will see when we focus on risk in Objective 1.10. The first thing you must do for business continuity planning (BCP) is to complete an inventory to understand and document what systems, information, equipment, facilities, and personnel support the critical business processes. This inventory is vital to complete a business impact analysis, discussed next.

Business Impact Analysis

A business impact analysis, or BIA, identifies the organization's critical business processes, as well as the systems, information, and other assets that support those processes. The goal is to determine which processes the business must absolutely maintain to carry out its mission and minimize financial consequences. A BIA helps prioritize assets for recovery should the organization lose them if it suffers an incident, such as a natural disaster, a major attack, or other catastrophe. The BIA directly informs risk management processes, as previously mentioned, because the inventory of business processes and supporting assets helps determine which security controls must be implemented in the infrastructure to protect those assets, thus lowering the risk of losing them.

Developing the BIA

Developing the BIA is a cooperative effort among the cybersecurity and IT personnel and business process owners. If the organization has not already developed a criticality list of its key business processes, and which processes it must keep up and running to function, then this is a good opportunity to do so. The key business process owners must develop documentation identifying their own mission and goals statements and how these business processes support the organization's overall mission and goals. They must inventory and list the key upstream processes that keep their specific business areas going. Then they must decompose those key processes down into the smallest level possible so that they understand the various relationships between processes, as well as dependencies involved, even with processes that may have previously been considered insignificant.

Once this process workflow is complete, cybersecurity analysts and IT personnel map the various systems and information flows that support those business processes to determine the impact of a loss at any point in the workflow. Often the BIA development project uncovers critical assets that no one previously thought were important, simply because they support other systems and dependencies. A BIA can map out the organization's entire business process infrastructure, as well as the critical assets that support those key areas.

Scope

The scope of the BIA should obviously cover the organization's critical business process areas, but first those processes must be discovered, formally documented, and prioritized for importance. Business process owners need to decide which processes are most critical and offer information on which processes are less important. They must then determine which processes are essential to maintain acceptable operations, which processes can afford to be down or nonfunctional for specific periods of time, and which processes are not critical but still necessary.

The scope of the BIA will depend on the impact values assigned to these key business process areas. In turn, the key information assets that support these critical business processes

must also be included in the scope of the BIA, once they are identified. They will then be prioritized in terms of maintainability and recoverability.

Documenting the BIA

The BIA should be formally documented in such a way that it is easy to understand and adequately covers the scope of what the organization is trying to accomplish. All key business process owners, as well as their staff, should be able to review the BIA and provide input and suggestions for improving it, and should ensure that nothing has been left out.

Once the BIA has been approved throughout the organization, everyone should be familiarized with it, and it should be stored in a secure area so that it cannot be easily altered. However, all authorized stakeholders should be able to access it to review it and make updates when needed. The BIA should be reviewed periodically to ensure that it is still current, especially with new changes in the infrastructure, new technologies, new risks, and so on.

Cross-Reference

Business continuity, along with disaster recovery, is discussed in much more detail in Domain 7.

REVIEW

Objective 1.8: Identify, analyze, and prioritize Business Continuity (BC) requirements In this objective we discussed the necessity for business continuity and the business impact analysis. Business continuity is concerned with keeping the critical business functions that support the mission maintained and operating, even during a major incident. A business impact analysis is a review process and the resulting document that determines what critical processes support the organization's mission, as well as the information assets that support those critical business processes. This includes systems, information, data flows, equipment, facilities, and even personnel. Business process owners take the first step in inventorying those critical processes, and then IT and cybersecurity personnel inventory and prioritize the assets that support them. The BIA must be appropriately socialized throughout the organization so everyone can have the opportunity to review it and propose changes, as well as know and understand its contents.

1.8 QUESTIONS

1. You are a cybersecurity analyst tasked with assisting in writing the organization's business impact analysis. Which of the following is the first step in writing the BIA?
   - **A.** Developing a disaster recovery plan
   - **B.** Performing a risk assessment
   - **C.** Inventorying all infrastructure assets
   - **D.** Documenting all critical business processes

2. You are developing a BIA and need to ensure that it is scoped correctly. Which of the following would *not* be part of the BIA scope?
   - **A.** Vulnerability assessment for all critical assets
   - **B.** Inventory of all critical business processes
   - **C.** Inventory of all information system assets that support critical business processes
   - **D.** Dependencies of the different business processes on various assets
3. Which of the following should take place after the business impact analysis process has been completed?
   - **A.** The BIA documentation should be secured away, with access restricted to senior managers due to its confidentiality.
   - **B.** The BIA documentation should be monitored for potential updates.
   - **C.** The BIA documentation should be submitted to an auditor for approval.
   - **D.** The BIA documentation should be included as part of the disaster recovery plan.

1.8 ANSWERS

1. **D** Identifying and documenting all critical business processes that support the organization's mission is the first step in preparing a BIA, since all other actions depend upon that determination.
2. **A** A vulnerability assessment is not part of the business impact analysis process scope. It is, however, critical to the overall risk assessment process.
3. **B** After the business impact analysis has been completed, it should be made available to all authorized stakeholders for periodic updates. The analysis should be monitored since business processes and supporting technologies sometimes change, which could affect the BIA. Submitting the BIA to an auditor for review is not required. A BIA is part of business continuity planning, not disaster recovery planning, which are two separate but related processes.



Objective 1.9 Contribute to and enforce personnel security policies and procedures

In Objective 1.7 we discussed policies, which are internal governance documents that support both external governance requirements (i.e., laws, regulations, and industry standards) and internal requirements set forth by management. Now we will look at the area of personnel security and associated policies under the administrative and management processes.

Personnel Security

The personnel security program is designed to identify controls to effectively manage the security-related aspects of hiring and retaining people in the organization. These activities include pre-employment practices and controls, on- and offboarding processes, termination processes, and personnel security training. For the most part, personnel security controls are administrative or managerial, but you will also occasionally find technical controls that fit into the personnel security function.

The personnel security program establishes good security practices, such as:

- Clearance/need to know
- Separation of duties
- Principle of least privilege
- Preventing collusion
- Ensuring that people are held accountable for their actions
- Preventing and dealing with insider threats
- Security awareness and training programs for employees

Cross-Reference

Security awareness and training programs are discussed in greater detail in Objective 1.13.

Candidate Screening and Hiring

Before potentially being hired, all candidates must be carefully vetted to ensure that they are the right fit for the organization. Most positions in modern companies, particularly in security, are considered positions of trust. Often, individuals must possess a minimum level of substantiated trust in the verifiable facts of their background, such as positive legal and educational standing, financial responsibility, staunch ethics, and so on. Substantiation is primarily done through a series of suitability checks, which could include

- Background checks, such as security clearance vetting
- Credit or financial responsibility checks
- Educational credential verification
- Review of criminal records

While a great deal of information will be generated from these different types of checks, the organization has to be cognizant of what information it cannot collect. Generally, information considered privacy related, such as past personal relationships, group or organization affiliations, political leanings, medical history, and so on, is considered off-limits as part of the screening process. Some of this information may be requested and provided by the employee later on in the process, such as relationship status for company insurance, but the organization should not request information that might be legally or ethically off-limits.

Employment Agreements and Policies

Onboarding employees should review the different policies that will affect their employment, such as acceptable use policies and equipment use policies. These policies may be included in a comprehensive employee package provided to the employee on their first day on the job. Employees must carefully review, understand, and attest via signature their understanding of, and agreement to, the policies that will be enforced as a condition of their employment. Examples of employee agreements and policies that may be required for an employee's review and signature on initial hiring include

- Acceptable use policy
- Equipment care policy
- Social media policy
- Data sensitivity/classification policy
- Harassment policy
- Safety policies
- Security incident reporting policies

Once an employee signs these policies, they become part of the employee's record and signify their pledge to comply. These policies are enforceable under law and employees can be disciplined or even terminated if they do not obey them. It's vitally important that an organization create and provide these policies for new employees so that they cannot later claim they had no knowledge of the policy or did not understand it. That's why it's important that the employer obtain the signature of the employee, signifying that they have read and understand the ramifications of the policies.



EXAM TIP Key personnel security policies that require special attention during the employee onboarding process include acceptable use, privacy, and data sensitivity. These policies may be all rolled into a single employee policy or be part of several other policies, but these key subjects should be addressed.

Onboarding, Transfers, and Termination Processes

The employment agreements and policies are key for the proper implementation of various employment processes. Ensuring that employees understand and accept (in written form) the organization's rules and expectations allows for consistent application of processes like onboarding and termination. Processes are also needed for actions that occur at various points during the lifetime of their employment, such as when they are transferred, promoted, or demoted. Regardless of the phase, these processes regularly require review and validation of an employee's access to sensitive systems and data.

Onboarding

The initial process of bringing an employee into the organization is referred to as *onboarding*. First, the employee must review and sign the employment agreements and policies previously described. Then, provisioning activities must be performed to provide the employee's username, password, rights, permissions, and privileges regarding systems and data.

Normally, the provisioning process is initiated as part of the hiring process and coordinated between the human resources department and the individual's business unit and supervisory chain. The business unit must determine what type of access the employee will have and to which resources. These permissions must be provided not only to HR but also to the IT department. Systems or data owners responsible for those assets may also be in a position to grant access. The IT department verifies the appropriate information about the employee from HR, such as security clearance, verifies the need-to-know for particular systems or data from the employee's supervisory chain, and creates the account granting appropriate access in the company's systems.

During this process, the employee is also issued any required equipment, such as a laptop, company smartphone, or security token. The employee should sign for this equipment and must agree to its care.

The employee will also receive any initial security awareness training during the onboarding process, which should reflect general security responsibilities, information about threats specific to the company or the employee's position, and so on.

Transfers, Promotions, and Disciplinary Activities

Personnel security doesn't stop with onboarding. There are activities that occur throughout the lifetime of the employee's service in the organization. Policies are frequently changed, and all employees must review them and sign again acknowledging their understanding and agreement to comply with them. Refresher security awareness training also happens periodically, as we will discuss later in Objective 1.13. There are also some particular events during the employee's career, such as transfers, promotions, and disciplinary events, that receive special handling with regard to personnel security.

Transfers between organizational departments or divisions receive special consideration because of the additional rights, permissions, privileges, and access to sensitive data or systems granted to the employee. These systems and data may have their own access requirements that must be modified when the employee is transferred. In addition to reviewing and signing additional policies or access requests, and being verified by the supervisory chain, employees may have to undergo specific training for more sensitive data or systems.

One issue with transfers is *privilege creep*, which can happen over time when an employee moves from job to job within an organization, or is simply promoted, and retains privileges they no longer need. Permissions should be reviewed for both the outgoing and incoming positions to ensure the employee only has the access needed and no more. This practice directly supports the principle of least privilege.

Demotions and disciplinary actions especially require privilege review, since these negative actions may necessitate that an employee be removed from specific programs or have their access restricted. Disciplinary actions should be recorded in the employee's records, including the reason and final adjudication of those actions. Management should monitor these employees more closely for a period of time to ensure that the demotion or disciplinary action does not trigger them to violate security requirements.

Terminations

Terminations can happen for a variety of reasons and do not always have to be negative in nature. Positive separations like a retirement or a routine job change may not be cause for any additional personnel security concerns. These types of terminations should follow a routine offboarding process where there is an orderly return of equipment, deactivation of accounts, return of sensitive data, reduction and elimination of access to sensitive systems, and an orderly departure from the organization.

Terminations for other than favorable reasons, such as violation of a policy, law, or firing for cause, may necessitate additional security measures if management is concerned about an individual destroying or stealing company property or endangering the safety of others. In such cases, once the decision has been made to terminate an individual, the organization must act swiftly and immediately revoke access to systems and data. The person should be escorted at all times within the organization, and there should be witnesses to any actions that the organization takes, such as the termination notification, security debriefings, equipment return, and so forth.

All onboarding, transfer, and termination procedures should be well documented and include information security considerations, such as provisioning and deprovisioning, data protection, nondisclosure agreements, as well as other HR documentation.

Vendor, Consultant, and Contractor Agreements and Controls

While not considered employees, organizations often have vendors, consultants, external contractors, and business partners working in their facilities. Although there may be unique security requirements to implement for these individuals, by and large most of the personnel security practices we have discussed apply to them. To access organizational systems, they must review and abide by policies such as acceptable use, equipment care, and so on. Access to sensitive systems and data will require the necessary approvals from the supervisory chain or data and system owners. Organizations may also require limited background checks of external personnel. Even a vending machine vendor who only comes into the facility once a week may warrant a limited background check simply to make sure they have no serious criminal background, since they will be allowed in the facility, often unescorted, in areas that are near sensitive work.

Most of the personnel security requirements imposed on external personnel are included in contractual agreements between the organization and their employer. By including the requirements in contracts, they are legally enforceable. If external personnel do not comply with the organization's security policies, they can be removed from the facility or contract, and the other company may incur liability from their actions. The key takeaways from this discussion are to ensure the personnel requirements are included in all contractual agreements and formalize a process to assure that any requested access is carefully vetted and documented.

Compliance Policy Requirements

Most policies are designed to ensure compliance with some type of governance requirement, whether that requirement is imposed by external governance, in the form of laws and regulations, or imposed by an organization's management to articulate their own requirements. In any event, implemented policies must be obeyed by employees. Many policies not only describe the requirements of the policy but also the consequences for noncompliance. These consequences could range from simple disciplinary measures, to suspension, or all the way to termination.

Some of the common personnel policies that could result in disciplinary actions or termination if not followed include the acceptable use policy, equipment care and use policies, harassment policies, safety policies, data protection policies, and social media policies. The bottom line here is that with all personnel security policies, there is a compliance piece to the policy that must be considered.

Privacy Policy Requirements

Privacy policies can be a double-edged sword that affect organizations in different ways. First, privacy policies must be implemented to protect both employee and customer personal data. Various laws and regulations have specific requirements for the privacy policy. The most prevalent laws or regulations include the U.S. Health Insurance Portability and Accountability Act (HIPAA) and the European Union's General Data Protection Regulation (GDPR). Both of these regulations determine what must be included in a privacy policy and how that policy must be implemented and enforced, such as:

- How and why personal data is collected from individuals, such as employees
- How that data will be used
- How the data will be stored or protected
- How the data will be disseminated to other entities
- How the data will be retained or destroyed when no longer needed

Second, in addition to protecting the data of individual employees and customers, privacy policies are also implemented to protect the organization. Often organizations are in possession of data that must be carefully protected. If this data were to be lost, stolen, or otherwise

compromised, the organization could be in legal trouble. Privacy policies, from the organization's perspective, often dictate how to protect sensitive personal data, such as healthcare or financial data. These policies help to fulfill due diligence and due care requirements for companies and demonstrate compliance with regulations.

REVIEW

Objective 1.9: Contribute to and enforce personnel security policies and procedures In this objective we discussed personnel security and focused on the different policies and processes organizations use to manage security of their personnel. Personnel security doesn't simply focus on employees, or managers; other personnel are included in those policies, such as vendors, consultants, and external contractors.

We discussed the policies and processes that go into initial candidate screening and hiring a candidate to make them a permanent employee. Employee agreements are necessary to ensure that new employees understand their rights and responsibilities and are an excellent way to initially inform new employees about their risk security responsibilities and then ensure they understand and agree to them.

Personnel activities, such as onboarding, employee transfers between organizational elements, and employee termination require strict adherence to security policies. These activities ensure that employees are indoctrinated properly, understand their security roles and responsibilities, and are managed throughout their tenure at the organization. Transfer procedures ensure that employees do not improperly accumulate privileges and that those privileges are examined and validated as employees change roles or job positions. Termination procedures ensure that there is an orderly transfer of knowledge, equipment, and data back to the organization when an employee has ended their relationship with the company. Effective termination processes help prevent equipment or data theft, avoid potential safety issues with personnel leaving the organization, and ensure the interests of both the employee and the company are considered.

External personnel that are essentially full-time employees, such as vendors, consultants, and contractors, are subject to certain personnel security policies, such as those that require security indoctrination and training, security clearances, need-to-know, background checks, and so on. These are put in place to ensure that personnel, even those that are not technically company employees, are made aware of their responsibilities and held accountable for their actions.

We also discussed compliance policies, which are certain policies that are created and enforced to maintain compliance with governance and directly affect the personnel that are part of an organization. Primarily focused on privacy and data protection, these policies detail the behavior and actions necessary to comply with internal and external governance requirements and describe consequences in the form of discipline or termination if they are not followed.

Privacy policies serve to protect the data of an organization, its customers, and its personnel. Privacy policies dictate how personal data is collected, used, stored, and disseminated. Privacy policies also serve to ensure compliance with external governance, such as laws and regulations.

1.9 QUESTIONS

1. Emilia is being vetted for employment in your organization. As part of the routine prescreening checks, the human resources department is running a background check on her. Which of the following is the most relevant piece of information for a position within your organization that requires Emilia be placed in position of trustworthiness?
   - **A.** Health history
   - **B.** Criminal record
   - **C.** Political leanings
   - **D.** Employment history
2. Evie is onboarding into the organization as a cybersecurity analyst in the threat modeling and research department. As part of her onboarding process, she must review and sign company policies that all employees are required to acknowledge. Additionally, because of her position, she must also be granted access to sensitive systems and data. Which of the following roles would determine and approve access to those sensitive systems?
   - **A.** Department supervisor
   - **B.** Human resources supervisor
   - **C.** Company president
   - **D.** IT security technician
3. Caleb is being transferred to a different department within the company and is receiving a promotion at the same time. His duties will be significantly different in the new department, and he will be supervising other personnel. Which of the following changes should be made to his access to sensitive systems and data?
   - **A.** He should continue to receive the privileges from his old department, and the privileges he needs for his new department should be added.
   - **B.** He should be carefully vetted for access to any new systems or data that come with his promotion and transfer, but his old permissions do not need to be reviewed.
   - **C.** His access to systems and data that are not required for his new position should be reviewed and removed, and he should be appropriately vetted for access to any new systems or data he requires as a result of his transfer and promotion.
   - **D.** He should immediately have his access to all systems and data in his old department removed and he should undergo a vetting to determine suitability for access to systems relevant to his new position.

4. Sam is an employee working in the accounting department. During routine auditing, it was discovered that Sam has committed fraud against the company. The decision has been made to terminate his employment. Which of the following should be completed as part of the termination process?
   - **A.** Immediately revoke his access to all systems and data.
   - **B.** Inform him that he has two weeks before he leaves the company and remove his access to systems and data on his last day of work.
   - **C.** Review his access to all systems and data and remove only the access to sensitive accounting systems.
   - **D.** Allow him to offboard the company unescorted and require him to turn in his equipment, data files, and access badges/tokens before he leaves for the day.

1.9 ANSWERS

1. **B** While employment history could be critical to determining experience and work history, a criminal record is a key piece of information in determining the suitability for trustworthiness in a sensitive position within the organization. Medical history and political leanings are irrelevant to a sensitive position in the organization, and neither type of information should ever be requested during the hiring process.
2. **A** The department supervisor should approve access to sensitive systems and data, as that person is likely the data or system owner and accountable for the security of those systems. Human resources cannot make any access determination since that is not their area of expertise. Access control decisions are normally delegated below the level of the company president, unless there are extreme or unusual circumstances. IT security personnel are normally responsible for provisioning accounts and access, not making access determinations.
3. **C** Caleb should have his access to any systems and data from his old department and position reviewed to determine which access he still requires, and access he no longer needs should be removed. He should be appropriately vetted for access to any new systems and data that come as a result of his transfer and promotion, assuming that was not part of the overall vetting process for those personnel actions.
4. **A** Since Sam is being terminated under other than favorable circumstances, such as the commission of fraud against the company, he should have his access to all systems and data terminated immediately. He should also be escorted throughout the facility, and his supervisor should accompany him as he turns in his equipment, data, access badges and tokens, and so on. All company personnel actions should be witnessed, such as debriefings, signing nondisclosure agreements, and so on.



Objective 1.10 Understand and apply risk management concepts

Risk is the probability (likelihood) that a threat (negative event), such as a disaster or malicious attack, will occur and impact one or more assets. Risk management is the overall program of framing, assessing, responding to, monitoring, and managing risk. In this objective we will cover the fundamental concepts of risk and risk management.

Risk Management

Risk management consists of all the activities carried out to reduce the overall risk to an organization. Although risk can never be completely eliminated, risk can be reduced or mitigated to a level that is satisfactory to an organization. To understand risk management, you must understand the elements of risk, as well as risk management processes and activities.

Elements of Risk

There are five general elements of risk that are considered within the cybersecurity community: an organization's assets, its vulnerabilities and threats, and the likelihood and impact of an event. Any number of external and internal factors can affect those components and, in turn, increase or decrease risk.

Assets

An *asset* is anything of value that the organization needs to fulfill its mission, such as systems, equipment, facilities, data and information, and people. Assets can be tangible or intangible. *Tangible assets* are items that we can easily see, touch, interact with and measure; examples are systems, equipment, facilities, people, and even information. Assigning a monetary value to tangible assets may be relatively easy, since we must consider replacement costs for systems and equipment, the cost of upgrading facilities, the revenue a system or a set of information generates, and how much we pay people in terms of labor hours. *Intangible assets* are those that cannot be easily interacted with or valued in terms of cost, revenue, or other monetary measurement, but are still critical to the organization's success. Intangible assets include items such as consumer confidence, public reputation, and prominence in the marketplace. These are all valuable assets that an organization must protect.

Vulnerabilities

Vulnerabilities can be defined in different ways. First, a vulnerability may be defined as a weakness inherent in an asset or the organization. For example, a system could have weak encryption algorithms built in that are easy to circumvent. Second, a vulnerability may be

defined as a deficiency in security measures or controls that protect assets, such as the lack of proper policies and procedures to secure assets.

Threats

A threat is a negative event that has the potential to exploit a vulnerability in an asset or the organization. Threats take advantage of weaknesses and attack those weaknesses, causing damage to an asset or the organization. A concept associated with threats that you need to understand for this objective is *threat actors* (also called *threat sources* or *threat agents*), which initiate or enable threats. Another important concept is that of threat and vulnerability pairing. Theoretically, threats do not exist if there is no vulnerability to exploit, and vice versa. Threats and vulnerabilities are often expressed together as a *threat-vulnerability pair*, even though some threats apply to more than one vulnerability.

Likelihood

The discussion of likelihood and impact is where we begin to truly define risk. Likelihood is often expressed as the probability that a negative event will occur—exploiting a vulnerability and causing damage (impact) to an asset or the organization. Likelihood can be expressed numerically, as a statistical number, or qualitatively, as a range of subjective values, such as very low, low, moderate, high, and very high likelihoods. Later in this objective we will discuss the methods of expressing likelihood and impact using these objective and subjective values. Likelihood can be determined using several methods, including historical or trend analysis of available data, probability and outcome, and even several subjective methods.

Impact

As mentioned earlier, impact is the level or magnitude of damage to an asset or even the entire organization if a negative event (the threat) occurs and exploits a weakness (a vulnerability) in an asset or the organization. As with likelihood, impact can be measured in various ways, including actual monetary loss if the asset is completely destroyed or requires extensive repairs, as a numerical percentage, or as a range of subjective values, such as very low, low, moderate, high, and very high impact.

Determining Risk

As stated above, risk is the probability (likelihood) that a threat (negative event), such as a disaster or malicious attack, will occur and impact one or more assets. Because the values of likelihood and impact vary, high risk could mean that the likelihood of a negative event is high or the level of impact is high. Since both elements function independently, even when the likelihood of an event occurring is low, if the potential damage to the asset is high, then the risk is high. Risk is often expressed in a pseudo-mathematical formula, Risk = Likelihood × Impact, which we will discuss later in the objective.

Identify Threats and Vulnerabilities

One of the key steps in risk management is identifying your assets. If you don't know what infrastructure is connected, how it exchanges data with other assets, and the importance of those assets, then you cannot manage risk. However, after you review and document assets, you must then identify the threats to those assets and the vulnerabilities that are inherent to them.

Identifying Threats and Threat Actors

As mentioned earlier, a threat is a negative event. Threats can be potential or realized. Once a threat has actually exploited a vulnerability, you need to determine the extent of damage to the asset or the organization. Fortunately, threats can be identified before they are realized and then matched to the vulnerabilities they would exploit in assets or the organization. Threats can be categorized in different ways—human-initiated or natural, intentional or accidental, generalized or very specific. Threats can target multiple vulnerabilities at once in an entire organization (think about a hurricane or flood) or they can target very specific vulnerabilities, such as a weak encryption algorithm in an operating system.

Threats can be identified generically by listing some of the common negative events that can affect an organization or its systems. There is a wide variety of threat libraries available to organizations from public sources that list threats and their threat sources. However, simply listing the threat does not allow the organization to discover exactly how a specific threat would affect it. A more effective process is called *threat modeling*, which looks at the organization and specifically matches likely threats with discovered vulnerabilities and organizational assets.

Cross-Reference

Threats and threat modeling are discussed in more detail in Objective 1.11.

Identifying Vulnerabilities

As mentioned earlier, a vulnerability is a weakness in an asset, or a deficiency in or lack of security controls protecting an asset. All assets have some sort of vulnerability, whether it is a vulnerability in the operating system that runs a server, the encryption algorithm that sends information across a network, an authentication method, or poorly written software code. But vulnerabilities are not tied simply to systems or data; vulnerabilities can exist throughout the administrative, technical, and physical processes of an organization. An organizational vulnerability might be a lack of policies or procedures used to secure its assets. Physical vulnerabilities may include an area around a facility where an intruder could easily enter the grounds.

Vulnerabilities are typically discovered during a process known as a *vulnerability assessment*. Vulnerability assessments are often technical in nature, such as scanning a system for configuration issues or lack of security patches. However, vulnerability assessments can also span other areas, such as administrative or business processes, facilities in the physical environment, and even vulnerabilities associated with human beings, such as those that might be present in a social

engineering attack. The other types of assessments that can expose vulnerabilities in an asset or the organization include risk assessments (discussed next), penetration tests, and even routine system tests. Vulnerabilities can be eliminated or reduced by implementing stronger security controls or correcting weaknesses in assets. We will discuss some of the methods of reducing risk associated with vulnerabilities later in the objective.

Risk Assessment/Analysis

In order for organizations to determine how much risk they can endure, they develop risk appetite and risk tolerance values. *Risk appetite* is a general term that applies to how much risk the organization is willing to accept. In risk-averse organizations, the risk appetite level is not very high. In organizations that allow and even encourage risk taking, in order to expand business, the risk appetite is higher.

Risk tolerance, on the other hand, typically applies to individual business ventures or efforts. *Risk tolerance* is essentially the variation or deviation from the risk appetite that an organization is willing to take, depending on how much the organization feels that variation is worth for that particular business effort. Risk tolerance could be slightly more than the organization's risk appetite for a given venture, or even somewhat less. These values for risk appetite and tolerance are developed from different factors, such as the organization's *risk culture* (how the organization as a whole feels about taking risk, such as being risk-averse), operating environment, governance, and many other factors.

These primary elements of risk, likelihood and impact, have to be determined before risk can be determined. In this two-step process, the risk assessment process happens first and consists of gathering data about the organization and its assets. The risk analysis process occurs afterward and involves looking at all the information the organization has gathered and determining how it fits together to define the risk to an asset or the organization.

Risk Assessment

The terms risk assessment and risk analysis are often used interchangeably; even some formalized risk frameworks, discussed a bit later, use them interchangeably. However, risk assessment and risk analysis are actually distinct and separate processes within the overall risk management program. A risk assessment often includes a risk analysis as part of its process. The overall risk assessment process involves gathering data and analyzing it to determine risk to the organization, assets, or both. The data collected is directly related to some of the elements of risk discussed earlier: assets, vulnerabilities, and threats. Likelihood and impact, the other two elements of risk, are generally calculated from that data during the analysis process.

The information required to determine risk can come from a wide variety of sources. Information about assets can come from inventories, network scans, business impact analysis, and so on. We also gather information about threats and vulnerabilities that affect those assets, through threat and vulnerability assessments. As mentioned previously, generalized information about threats is easily obtained but does not offer a level of depth or detail useful

in determining how likely it is that a given threat will attempt to exploit a specific vulnerability in an asset. Again, this is where threat modeling comes in, which we will discuss in Objective 1.11.

Several risk frameworks prescribe detailed risk assessment processes. For example, the National Institute of Standards and Technology (NIST) Risk Management Framework (RMF) details a four-step risk assessment process in its Special Publication 800-30 (currently Revision 1):

1. Prepare for the assessment.
2. Conduct the assessment:
   a. Identify threat sources and events.
   b. Identify vulnerabilities and predisposing conditions.
   c. Determine likelihood of occurrence.
   d. Determine the magnitude and impact.
   e. Determine risk.
3. Communicate results.
4. Maintain the assessment.

In this example, the risk analysis portion falls under step 2, conduct the assessment. In addition to identifying information about threats and vulnerabilities, it also involves determining the likelihood of a negative event occurring, as well as estimating the impact to the asset or the organization. We will go into a bit more depth on risk analysis next.

Risk Analysis

Risk analysis occurs after gathering all the available data on assets, threats, and vulnerabilities. In addition to these elements of risk, information on various risk factors—a variety of elements that can affect risk in the organization—is also gathered. Risk factors are things the organization may or may not be able to control that influence some of the risk elements, such as the economy, the organization's standing in the marketplace, the internal organizational structure, governance, and so forth. For example, the economy can affect the value of an asset, how much revenue it brings in, and the cost to repair or replace the asset. Governance can affect the level and depth of security controls that must be present to protect a given type of data. Internal organizational structure can affect who owns business processes and how much resources are committed to them. Information on these risk factors is included in the "predisposing conditions" portion of gathering information during the assessment.

The purpose of risk analysis is to determine the last two elements of risk: likelihood and impact. Likelihood, as previously noted, considers the following factors:

- The probability that a threat event will occur
- That it will be successful in exploiting a given vulnerability, and
- That it will cause some level of damage to an asset or the organization.

As discussed earlier, likelihood can be expressed in terms of statistical percentages or in subjective terms, and impact typically is expressed in numerical values but also can be expressed subjectively.

Risk analysis uses two primary means to qualify risk: quantitative and qualitative. *Quantitative analysis* focuses on concrete, measurable data, usually in the form of numbers. For example, using historical analysis and statistics, we can calculate the expected degradation an asset may experience during specific events. We may expect to lose 25 percent of an asset during a flood, so multiplying that number (called the *exposure factor*) by the asset value will tell us how much, in numerical terms, we could expect to lose. Other types of numerical calculations involve data such as time, distance, financials, statistical calculations, and so on. The point is that quantitative analysis is objective and uses hard data.

Qualitative analysis, on the other hand, uses subjective data. This means that the data is based upon opinion or is subjective in nature. Instead of assigning numerical values, we may assign qualitative values, as previously mentioned, such as very low, low, moderate, high, and very high. The problem with these values is that what may be considered a "very high" value to one person may simply be a "high" value to someone else. Plus, the quality of the data used to make the calculations may be subjective and opinion-based. This does not mean that qualitative analysis is an inferior form of analysis; in fact, most risk analysis uses qualitative judgment, since many intangible assets are difficult to quantify. More often than not, you'll find that risk analysis uses a combination of quantitative and qualitative methods. When referring to impact or loss, numerical data is most often used because financial impact is real and measurable, and quite meaningful to senior executives in the company. Likelihood is naturally more subjective in nature since it may be based on extrapolation of data.

One of the quantitative measurements you should understand for the CISSP exam is how to calculate the single loss and annualized loss of an asset due to the manifestation of a threat, which is directly related to the impact element of risk. Table 1.10-1 lists some of the more common quantitative formulas you should remember for the exam.

To use these formulas, you need a few critical pieces of information:

- Asset value (AV) is the calculated value of how much the asset is worth, in terms of cost to replace, original purchase price, amount of revenue the asset generates, or some other monetary value the organization places on the asset.

TABLE 1.10-1 Quantitative Risk Formulas

To calculate:	Use this formula:	Description
Single loss expectancy (SLE)	SLE = AV × EF	Calculates a single loss of an asset, based on the asset's value and how much of the asset would be lost in a given event.
Annualized loss expectancy (ALE)	ALE = SLE × ARO	Calculates how many losses for the given asset are expected per annum.

- Exposure factor (EF) is expressed as a percentage, and is the portion of the asset that can be expected to be lost during a negative event.
- Single loss expectancy (SLE) is expressed in monetary terms and represents the dollar value of the loss due to a single negative event.
- Annualized rate of occurrence (ARO) is a value that expresses how many times per year the event can be expected to occur, resulting in a loss. This value is 1 for once a year, .5 for every two years, and so on. If the event that causes a loss occurs more than once a year, the number will be greater than 1.
- Annualized loss expectancy (ALE) is the amount of loss expected for a given asset due to a specific threat event on an annual basis.

Note that these are very simplistic formulas as they only account for a single event with a single asset. You need to aggregate multiple events, determine the value of many different assets, and then roll up the results for a more complete picture of risk, which is why quantitative analysis is rarely performed alone. Qualitative analysis is better suited to roll up risk from a single asset to the entire organization, given multiple threat events, assets, and various other factors.



EXAM TIP Understand the differences between quantitative analysis and qualitative analysis and be familiar with the formulas for SLE and ALE for the exam.

Risk Response

Risk response is what an organization does after it has thoroughly analyzed its risk and identified the actions required to reduce or mitigate the risk. Risk response seeks to lower the likelihood and impact of risk. If either of these two elements is reduced, then overall risk is reduced. Note that you can mitigate or even completely eliminate vulnerabilities, but you cannot eliminate a threat actor or threat event—you can only increase your defenses against it.

There are four general approaches an organization can take to manage risk:

- Risk mitigation (reduction)
- Risk transfer (or risk sharing)
- Risk avoidance
- Risk acceptance

Risk mitigation involves lowering risk by reducing likelihood or impact, often by eliminating or minimizing vulnerabilities or strengthening security controls. The goal is to reduce the level of total risk.

Risk transfer requires the offloading of some risk to another entity. A prime example of risk sharing or transfer is the use of insurance. It lowers the financial impact to the organization

should a serious negative event occur. Note that risk transfer is not meant to take away responsibility or accountability from an organization; the organization must still bear both of these, but it is not as likely to be impacted financially. Another example of risk sharing is the use of third-party service providers, such as those that may provide cloud services, hosted infrastructure, or even security services.

Risk avoidance does not mean that the organization simply turns a blind eye to risk. It means that the organization will avoid or cease performing activities that incur an unacceptable level of risk. The organization avoids activities, such as a new business venture, that may be beyond its risk appetite or risk tolerance levels.

Risk acceptance doesn't mean that the organization simply accepts the risk as is. It uses the other available responses as much as possible to reduce, transfer, or avoid risk, and whatever risk remains (called *residual risk*) is accepted if it is within risk appetite or tolerance levels.

Risk Frameworks

Risk frameworks provide a formal, overarching set of processes and methodologies that an organization can use to establish and run its risk management program. Some of these frameworks are driven by the organization's market or industry; other frameworks are promulgated by private organizations; and still others are published by government agencies. Most risk frameworks provide a structure within which to frame risk (determine the organization's risk appetite and tolerance levels), assess risk, respond to it, and monitor it. Some popular examples of risk frameworks include

- NIST Risk Management Framework (previously introduced)
- ISO/IEC 27005
- Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE)
- Factor Analysis of Information Risk (FAIR)

Countermeasure Selection and Implementation

Countermeasures (security controls) are used to reduce risk and are selected based on different factors. Most important is the ability to reduce risk, but there are other considerations like cost. If the use of the countermeasure costs more than the asset would cost to be repaired or replaced, then implementing that countermeasure may not be cost-effective. Countermeasures must be selected based on a cost/benefit analysis. The benefit of using the control or countermeasure must outweigh the cost of implementing and maintaining the control, as well as be balanced with the cost or monetary value of the asset itself. If the asset costs less to repair or replace than the countermeasure, is it worth putting in place the countermeasure to prevent the threat from damaging the asset? Remember that the value of an asset is not only its repair or replacement cost; the value should also consider the amount of revenue the asset generates and its overall value to the business processes.



EXAM TIP Although the terms control and countermeasure are almost synonymous, there is a subtle distinction: a control typically means an ongoing security mechanism to prevent a negative result, such as a compromise of confidentiality, integrity, or availability. Technically, a countermeasure is applied as a response after a compromise has occurred, such as during a malicious incident. Controls are preventative, whereas countermeasures are reactive. As the CISSP exam objectives frequently use these terms interchangeably and synonymously, we will also do so in this book.

Applicable Types of Controls

There are generally three types of controls that security practitioners value, as well as six control functions. Most controls are categorized as only one type of control but could be used for more than one function, depending on the context. The following are the definitions of the control types and functions you will encounter on the exam.

Control Types

The major control types are administrative (also referred to as *managerial*), technical (or *logical*) controls, and physical (or *operational)* controls. Table 1.10-2 describes these control types.

Control Functions

Control functions describe what a control does. While most controls are classified into one control type, controls can span more than one function. There are generally six control functions that you should remember for the exam, as listed in Table 1.10-3.

Note that controls can span multiple functions; for example, a video camera placed in a strategic spot can deter someone from committing a malicious act and it can also detect if a malicious act is committed. Deterrent controls must be known by an individual in order to

TABLE 1.10-2 Control Types

Control Type	Description	Example
Administrative (managerial)	Controls imposed by the organization's management	Policies and procedures
Technical (logical)	Controls using hardware and software	Firewalls, encryption mechanisms, authentication mechanisms, etc.
Physical (operational)	Controls pertaining to the physical environment	Gates, guards, fencing, physical alarm systems, temperature and humidity controls

TABLE 1.10-3 Control Functions

Control Function	Description	Examples
Deterrent	Deters an individual from committing a malicious act or violation of policy	Visible video cameras, signage, physical obstructions, computer warning banners
Preventive	Prevents an individual from committing a malicious act or violation of policy	Firewall rules, locked facilities, guards, object permissions
Detective	Detects a violation of policy or malicious act	Audit logs, intrusion detection systems, physical alarm systems
Corrective	Temporary measure that corrects an immediate security issue	Guards, fencing, rerouting network traffic in the event of an attack.
Compensating	Longer-term measure employed when a preferred control cannot be implemented	Additional security devices, stronger encryption methods, physical obstacles and barriers
Recovery	Controls used to bring a damaged or compromised asset back to its original operational state after an incident or disaster	Data backups, redundant spares

deter them from committing a malicious act or violating a policy; however, a preventive control does not have to be known in order to work. Additionally, a deterrent control is not always effective if the individual simply chooses to commit the act, while a preventative control will definitively help stop the act from being committed.

Another distinction to make is between corrective and compensating controls. A corrective control is temporary in nature and only serves to fix an immediate security issue. A compensating control is longer-term and may be employed when the organization can't afford a primary or desired control.

Control Assessments (Security and Privacy)

Controls must be periodically assessed for effectiveness, compliance, and risk.

- **Effectiveness** How effective are the security controls at protecting assets?
- **Compliance** Are the security controls compliant with required governance?
- **Risk** How well do the security controls as implemented reduce or mitigate risk?

When assessing a control, the organization wants to see how well the control is doing its job in protecting assets or, in the case of privacy controls, how well the control is protecting

individual data and conforming to the privacy policies of the organization. Controls can be tested in four main ways:

- Interviews with key personnel (system administrators, engineers, privacy practitioners, etc.)
- Observing the control in operation (determining if the control is doing what it is supposed to do)
- Documentation reviews (design and architecture documents, logs, maintenance records, etc.)
- Technical testing (e.g., system vulnerability scanning or penetration testing)

Controls should be tested on a periodic basis, and may be tested through specific control assessments, vulnerability assessments, risk assessments, system testing, or even penetration testing. Controls that fail any test for effectiveness, compliance, or risk reduction should be evaluated for replacement, upgrade, or strengthening. Results of control assessments must be thoroughly documented in an appropriate report and become part of the organizational risk posture.

Monitoring and Measurement

All elements of risk should be monitored on a continual basis, since threats change, new vulnerabilities are continually discovered, and even likelihood and impact can change, depending upon the organizational security posture and its operating environment. Risk is not static and must be monitored for any changes, which should, in turn, cause the organization to change its responses to meet any new or increased risk.

Risk, as well as its individual components, should also be measured. We briefly mentioned quantitative and qualitative measurement, each with its advantages regarding the type of data it requires and how meaningful it may be to the organization.

Reporting

Risk reporting is normally a formal process, based on the requirements of the organization or any governing entities that may require specific reporting procedures for compliance purposes. Since most types of security assessments fall under the overarching umbrella of risk assessments, the results of these assessments are reported as they are completed, so formalized risk reporting occurs on a fairly regular basis in most mature organizations. A key part of the formalized risk reporting process is what's known as the *risk register*, or sometimes known as a Plan of Action and Milestones (POA&M). Both documents record a variety of data, including risks, the assets they affect, the vulnerabilities that are part of those risks, and a plan for mitigating or responding to those risks . They may also assign risk owners and a timeline for addressing risk.

Informal risk reporting also happens as vulnerabilities are discovered or when risk factors affecting threats, vulnerabilities, impact, or likelihood are encountered. These risk factors could be things such as a security budget decrease or a new law or regulation that applies to the organization. Since these affect risk, they are often reported informally or may be recorded later in a formalized report.

Continuous Improvement

In the context of risk management, continuous improvement means that the organization must continually strive to assess, reduce, and monitor risk. This means continually improving its security processes, but also improving its security posture so that assets are better protected from threats and vulnerabilities. In a highly mature organization, a concept called *risk maturity modeling* may take place. Maturity models are designed to help organizations determine how well they perform their management activities. Maturity models are usually expressed in terms of levels (e.g., 1–5) that may show the organization is performing risk management in an ad hoc, unmanaged manner; in a repeatable manner where most procedures are documented and followed; or even all the way to a level where risk management processes are ahead of the game and proactively seek to manage risk based on data and predictive models.

REVIEW

Objective 1.10: Understand and apply risk management concepts In this objective we looked at risk management. We discussed the elements of risk, which consist of assets, vulnerabilities, threats, likelihood, and impact. Risk is a combined measure of the latter two elements, likelihood and impact. We also discussed how to identify threats and vulnerabilities. Threats are events that can exploit a vulnerability (a weakness) in an asset. Risk assessments consist of gathering data regarding assets, threats, and vulnerabilities and analyzing that data to produce likelihood and impact values, which make up risk.

We also discussed four risk response actions an organization can take: risk reduction or mitigation, risk transference or sharing, risk avoidance, and risk acceptance. We further listed a few risk frameworks that you may encounter during your risk management activities, such as the NIST RMF and ISO/IEC 27005. We also addressed countermeasure selection, which involves a cost/benefit analysis based on how much risk the control or countermeasure mitigates versus how much the control costs to implement and maintain. This must be balanced with the value of the asset. Controls that cost more to implement and maintain than the asset is worth may not be cost-effective.

We also examined types and functions of controls. There are normally three types of controls—administrative, technical, and physical. There are six control functions: deterrent, preventive, detective, corrective, compensating, and recovery. We discussed how to perform control assessments, which include assessing controls for effectiveness, compliance, and risk. The four ways to conduct control assessments are to interview key personnel, review documentation related to the control, observe the control in action, and perform technical testing on the control. We also briefly mentioned how risk and controls are

monitored and measured on a continual basis and how you should report risk and control results. Continuous improvement means that we must always strive to improve our security processes, controls, and risk management activities.

1.10 QUESTIONS

1. You are performing a risk assessment for your company and gathering information related to a lack of or inadequate controls protecting your assets. Which of the following describes this lack of adequate controls?
   - **A.** Threats
   - **B.** Vulnerabilities
   - **C.** Risk factors
   - **D.** Impact
2. You are performing a risk analysis but have found it is difficult to assign numerical values to some of the data collected during the analysis. You want to be able to express, using your expertise and fact-based opinion, values regarding the severity of risks to the organization's assets. Which of the following describes the method you should use?
   - **A.** Statistical
   - **B.** Quantitative
   - **C.** Qualitative
   - **D.** Numerical
3. Which of the following is the most important factor in selecting a control or countermeasure?
   - **A.** Cost
   - **B.** Level of risk reduction
   - **C.** Ease of implementation
   - **D.** Complexity
4. Which of the following is an effective means of formally reporting and tracking risk?
   - **A.** Risk register
   - **B.** Quantitative analysis
   - **C.** Risk assessments
   - **D.** Vulnerability assessments

1.10 ANSWERS

1. **B** A vulnerability is either a weakness in an asset or the lack of or inadequate controls protecting that asset.
2. **C** Qualitative analysis enables the expression of values for data that is difficult to quantify; these values are subjective and are fact-based opinion. All the other options describe quantitative analysis.

3. **B** The most important factor in control or countermeasure selection is the amount of risk it reduces; however, a cost/benefit analysis must be performed to determine if the amount of risk that is reduced balanced with the cost of the control to implement and maintain exceeds the value of the asset.
4. **A** A risk register is a tool used to formally report and track risks in the organization, and includes risk findings, vulnerabilities, the assets that incur the risks, and risk ownership, along with a timeline for mitigating risk and resources that must be committed to it.



Objective 1.11 Understand and apply threat modeling concepts and methodologies

In Objective 1.10 we discussed what a threat is and its associated terms, such as threat actor, threat source, and so on. You learned that a threat is a negative event that has the potential to exploit a vulnerability in an asset or the organization. In this objective we're going to look at various aspects of threats, including threat modeling, threat components, threat characteristics, and threat actors. While these aspects are sometimes discussed separately, they are all interrelated and contribute to each other.

Threat Modeling

Simply identifying a broad range of threats gives you a general idea of the things that can harm the organization; however, simply identifying generalized threats that may or may not affect your organization does not go very far in helping you focus on the particular threats that are targeting your specific assets. That's where *threat modeling* comes in. It involves looking at which specific threats are targeting your organization and assessing the likelihood that they will actually attempt to exploit a specific vulnerability in an asset and whether they could be successful in that exploitation. Threat modeling looks at all of the generalized threats and attempts to narrow them down based on realistic parameters, such as the assets you have, why someone or something would target them, and what realistic vulnerabilities might be present that they could exploit.

Cross-Reference

Threat modeling is also discussed in Objectives 1.10, 3.1, and 7.2.

Threat Components

Threats have many different facets and can be characterized in a variety of ways. Threats are made up of varying properties, including the source of the threat, the characteristics of the threat, whether the threat is potential or realized, and the vulnerability it can exploit.

Some threats target very specific vulnerabilities and may therefore be easier to manage; some threats, such as natural disasters, are more general and can wreak havoc across a wide variety of vulnerabilities and assets. Let's discuss a few threat properties.

Threat Characteristics

Before they occur, they are merely potential threats, but once they're actually initiated and take place, they are considered threat events. Remember that threat events must always exploit a vulnerability, which is why we typically see threats and vulnerabilities paired together. A threat event can be the destruction of data during a natural disaster, for example, or the actual exploitation of a vulnerability through malicious code. We often have small pieces of data that by themselves are meaningless, but when put together show that a threat has actually materialized and exploited a vulnerability. These are called *threat indicators*. When they are viewed collectively and show that a malicious event has taken place, they are called *indicators of compromise*.

Threats can be characterized in different ways, including

- Natural versus human-initiated (e.g., hurricanes versus hacking attempts)
- Potential versus actual (i.e., threats that have not occurred yet versus threats that have taken place)
- Threat source (e.g., hacker, complacent user, thunderstorm, etc.)
- Generalized versus specific (e.g., general threats against unpatched operating systems versus a threat that has been specifically identified to exploit a particular vulnerability)
- Known versus unknown (identified and categorized threats versus zero-day exploits, for example)

Threat Actors

Threat actors, also referred to as *threat agents* or *threat sources*, are entities that initiate a threat, promulgate a threat, or enable a threat to take place. Threat actors are not always human, although we ascribe most malicious acts to human beings. There are natural threat sources as well, such as floods, hurricanes, and tornadoes. Remember that threat sources can also be classified in different ways, as we mentioned also in Objective 1.10:

- **Adversarial** Malicious entities such as individuals, groups, organizations, and even nation-states
- **Accidental** Complacent users
- **Structural** Equipment or software failure
- **Environmental** Natural or human-initiated disasters and outages



EXAM TIP Although sometimes difficult to do, remember that you must try to differentiate the threat actor or source from the threat itself. Sometimes these are almost one and the same, but given enough information and context, you can distinguish the source of a threat from the threat, which is the event that occurs. Remember that a threat also exploits a vulnerability; threat actors do not. They merely initiate or enable the threat.

Threat Modeling Methodologies

Various formalized methodologies have been developed to address the different characteristics and components of threats. Some of these methodologies address threat indicators, some address attack methods that threat sources can use against organizations (called *threat vectors*), and some allow for in-depth threat modeling and analysis. All these methodologies allow the organization to formally manage threats and are critical components of the threat modeling process. A few examples are listed in Table 1.11-1.

While an in-depth discussion on any of these threat methodologies is beyond the scope of this book, research them and ensure you have basic knowledge about them for the exam.

TABLE 1.11-1 Various Threat Modeling Methodologies

Threat Model	Description
MITRE ATT&CK Framework	Public knowledge database of threat tactics and techniques
Diamond Model of Intrusion Analysis	Analytical model used to view the characteristics of threat actors/events and assists in the analysis to defend against them
Cyber Kill Chain	Cybersecurity model originally developed by Lockheed Martin to identify the various stages of threats during a cyberattack
OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation)	Methodology developed by Carnegie Mellon which focuses on operational risk, security controls, and security technologies in an organization
Trike	Open source threat modeling methodology focused on auditing
STRIDE (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege)	Threat modeling methodology created by Microsoft for incorporating security into application development
VAST (Visual, Agile, and Simple Threat modeling)	Threat modeling methodology incorporated into the software development life cycle (SDLC) and frequently used in Agile development models
PASTA (Process for Attack Simulation and Threat Analysis)	Threat modeling methodology focused on integrating technical requirements with business process objectives

REVIEW

Objective 1.11: Understand and apply threat modeling concepts and methodologies In Objective 1.11 we discussed the basic concepts of threat modeling. Threat modeling goes beyond simply listing generic threats that could be applicable to any organization; threat modeling takes a more in-depth, detailed look at how specific threats may affect an organization's assets and vulnerabilities. Threat actors include those that are adversarial and non-adversarial, such as humans and natural events, respectively. Various threat modeling methodologies exist to assist in this effort, including STRIDE, VAST, PASTA, and many others.

1.11 QUESTIONS

1. You are a member of the company's incident response team. Your company has just suffered a malicious attack, and several key hard drives containing critical data in various servers have been completely wiped. The initial investigation indicates that a hacker infiltrated the infrastructure and ran a script to delete the contents of those critical hard drives. Which of the following statements is correct regarding the threat actor and threat event?
   - **A.** The hacker is the threat actor, and the data deletion is the threat event.
   - **B.** The script is the threat actor, and the hacker is the threat event.
   - **C.** The script is both the threat actor and the threat event.
   - **D.** The data deletion is the threat event, and the script is the threat actor.
2. Nichole is a cybersecurity analyst who works for O'Brien Enterprises, a small cybersecurity firm. She is recommending various threat methodologies to one of her customers, who wants to develop customized applications for Microsoft Windows. Her customer would like to incorporate a threat modeling methodology to help them with secure code development. Which of the following should Nichole recommend to her customer?
   - **A.** PASTA
   - **B.** Trike
   - **C.** VAST
   - **D.** STRIDE

1.11 ANSWERS

1. **A** In this scenario, the hacker initiates the threat event, and the actual event is the data deletion from the critical hard drives. The script may be a tool of the attack, but it neither initiates the threat nor is the threat itself, since by itself a script doesn't do anything malicious. The negative event is the data deletion.

2. **D** STRIDE (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege) is a threat modeling methodology created by Microsoft for incorporating security into application development. None of the other methodologies listed are specific to application development, except for VAST (Visual, Agile, and Simple Threat Modeling), but it is not specific to Windows application development.



Objective 1.12 Apply Supply Chain Risk Management (SCRM) concepts

A company's suppliers introduce serious security considerations into the organization, particularly since supplies, including software and hardware, may be compromised with malware, be substandard in function or performance, or even be counterfeit. In this objective we discuss how an organization could reduce potential security issues involved with the supply chain.

Supply Chain Risk Management

The supply chain organization consists of the sequence of suppliers of hardware, software, and goods and services. An organization can have an upstream supply chain that provides its supplies, and a downstream supply chain that it uses to supply others. Supply chain risk management (SCRM) consists of the measures an organization takes to ensure that every link in the supply chain is secure. Many of the issues associated with SCRM include faulty or counterfeit parts, components with embedded malware, and other malicious acts. This objective looks at the risks involved in supply chain management and how the organization should address them.

Risks Associated with Hardware, Software, and Services

Any link in the supply chain can be attacked, but three common targets are hardware, software, and services. As with the different risks associated which each of these targets, there are different steps an organization can take to mitigate or reduce those risks. The next few sections describe these risks and measures.

Hardware

Electronic components and machine parts are often the target of attackers in the supply chain. Hardware risks include

- Faulty components or parts that do not meet the specified standard
- Counterfeit or fake parts passed off as the real part
- Electronic components loaded with firmware-level malware

These risks could lead to failures in critical systems due to substandard parts; legal ramifications because of counterfeit or fake parts bought and sold; and malware that may eavesdrop or steal information from electronic systems and send that information back to a malicious third party. These risks can be addressed using several methods, including vendor due diligence in checking the source of parts and tracking their interactions with other entities along the supply chain, third-party verification of hardware, and testing of parts prior to acceptance or use in critical systems.

Software

Software can present the same risks as hardware, including embedded malware or other suspicious code that may not perform to the standards the organization requires; faulty code that may not meet performance or function requirements; and counterfeit or pirated software, which may get the organization into trouble from a legal perspective. The methods used to combat software issues in the supply chain are almost the same as those used to combat hardware issues. The organization should use due diligence to ensure software is acquired from reputable vendors, who have solid, secure development methodologies; perform extensive software testing prior to acquisition or implementation of the software in critical systems; and seek third-party verification and certification of the software.

Services

Although often not considered in the same realm as hardware and software, services offered through the supply chain can also be subject to attack and compromise. Consider services that are often contracted out to a third party, such as security, e-mail, directory services, infrastructure, and software programming, all of which are subject to attack and compromise. Organizations could suffer from faulty or compromised software, services that are below the level of performance and function expected in the contract, and even malicious insiders within the third-party provider (consider data theft).

Organizations have a few methods to reduce third-party service provider risk, which include

- Ensuring the service level agreement (SLA) or contract includes clear, delineated security roles and responsibilities for both the service provider and the organization
- Reviewing the security program of the service provider
- Conducting audits on the service provider either by the organization or a third-party assessor

- Reviewing the service provider's own security assessments, if available
- Conducting the organization's own tests and security reviews for any services provided to the organization by the service provider
- Implementing nondisclosure agreements (NDAs)
- Understanding the legal and ethical environment in which the service provider operates



EXAM TIP Comprehensive agreements with service providers and suppliers are the key to lowering supply chain risk to acceptable levels; ensure the terms of the agreement include items such as expected service level, nondisclosure, security reviews and assessments, and legal liability.

Third-Party Assessment and Monitoring

There are two different contexts that address the topic of third-party assessment and monitoring:

- Assessing and monitoring a third-party service provider that provides services to your organization
- Using a third party to assess or monitor a service provider that provides services to your organization

In the first case, the organization should take steps to make sure that the authorization to assess or monitor the service provider is included in the SLA or contract. If not, the organization may not have the legal standing to do so. The organization should include any requirements levied on the provider regarding security assessments during the system or software development life cycle for any hardware or software provided by the third party. The organization should also have the ability to review those test results and provide input if the software or hardware does not meet the organization's required security specifications. The organization should also have the ability to call in a third-party assessor in the event laws or regulations require an independent assessment.

In the second case, bringing in a third-party assessor to review the performance of a provider is not a consideration to be taken lightly, although it may be required by law, regulation, or the industry governance. Finding a qualified third-party assessor can be expensive. For instance, payment card industry assessors must be certified and qualified by an independent organization to perform PCI DSS security assessments on organizations that manage credit card transactions. Industry standards often require these assessments periodically, so the third-party service provider may be under one of those requirements, which often requires them to foot the bill for the assessment, rather than the organization.

Cross-Reference

Objective 1.4 discussed PCI DSS security requirements more in detail.

Minimum Security Requirements

When engaging a third-party service provider, the organization should ensure that security standards and requirements are included in the language of the contract or service level agreement. Documented security requirements are especially critical for industries with regulatory requirements, such as the healthcare industry and the credit card industry, which are required to comply with the Health Insurance Portability and Accountability Act (HIPAA) security standards and the Payment Card Industry Data Security Standards (PCI DSS), respectively.

Even if no regulatory standards are imposed on a third-party provider, the organization has the ability to include and enforce the standards in any contract documentation. At minimum, the requirements should include specifications for access control, auditing and accountability, configuration management, secure software development, system security, physical security, and personnel security. Rather than draft its own standards, the organization could impose industry standards such as the National Institute of Standards and Technology (NIST) Special Publication 800-53 controls, the Center for Internet Security (CIS) Controls, or the ISO/IEC 27001 framework.

Cross-Reference

Table 1.3-2 in Objective 1.3 described these frameworks in more detail.

Service Level Requirements

An organization can impose service level requirements not only on its supply chain and third-party service providers, but also on suppliers of critical parts and software. Software and hardware suppliers must be held to a minimum standard of on-time delivery, quality, and delivery under secure conditions. Third-party service providers should meet benchmarks that include minimum uptime, incident response, hardware or software failures, and so on. These are all typically included in the contract, but also should be included in the previously mentioned document called the *service level agreement*, or SLA. The SLA is key to imposing and enforcing requirements on any links in the supply chain, regardless of whether it is hardware, software, or services.

REVIEW

Objective 1.12: Apply Supply Chain Risk Management (SCRM) concepts In this objective we discussed the basics of supply chain risk management, including the definition of supply chain, upstream and downstream suppliers, and risks associated with three key pieces of the supply chain. We discussed risks associated with hardware, which include faulty, compromised, or counterfeit hardware, and risks associated with software, which also include faulty, compromised, or even counterfeit software. We covered the third piece of the supply chain, which is services that may be contracted out to third-party providers.

Risks inherent to services include lack of security controls, malicious insiders, and faulty security processes. We also talked about third-party monitoring and assessment, both for the party providing services and the use of an external assessor. Finally, we discussed minimum security requirements that should be imposed on any type of service provider, or anyone else in supply chain, as well as the importance of service level agreements in imposing and enforcing security requirements.

1.12 QUESTIONS

1. Your company has decided to include supply chain risk into its overall risk management program. Your supervisor has tasked you with starting the process. Which of the following should you do first to begin supply chain risk management?
   - **A.** Conduct a risk analysis on the supply chain.
   - **B.** Identify all the upstream and downstream components of the company's supply chain.
   - **C.** Begin checking any received hardware for faults or compromise.
   - **D.** Begin scanning any purchased software for vulnerabilities.
2. Your company receives both hardware and software components from various overseas suppliers. As part of your effort to gain visibility on your supply chain risk, you decide that your company must start verifying hardware and software components more carefully. Which of the following is the best way to accomplish this?
   - **A.** Perform security testing on any hardware or software components received.
   - **B.** Request security documentation on any hardware or software components from the supplier.
   - **C.** Install hardware and software into critical systems and then test the systems.
   - **D.** Contract a third-party assessor to assess and monitor your suppliers.
3. Your company contracts infrastructure services from a local cloud service provider. When the contract was first written, security considerations were not included in the agreement. Now the contract is being renegotiated at the end of its term and your supervisor wants you to include several key requirements in the new contract. Which of the following should be included as part of the security requirements in the new contract?
   - **A.** NIST or CIS control standards
   - **B.** Incident response team readiness
   - **C.** Minimum security requirements to include controls and security responsibilities
   - **D.** Data confidentiality requirements

4. Your company processes sensitive data, and some of it is under regulatory requirements. You are contracting with a new third-party provider who will have access to the sensitive data. While regulatory requirements for protection of sensitive data will automatically be imposed on the new provider, which of the following should you also have in place to help protect sensitive data when it is accessed by the provider's personnel?
   - **A.** Nondisclosure agreement (NDA)
   - **B.** Service level agreement (SLA)
   - **C.** Provider's own internal security assessment report
   - **D.** Third-party assessor report on the provider

1.12 ANSWERS

1. **B** Before you can do anything else, you should take the time to identify all upstream suppliers for the company and the goods and services they provide, as well as any downstream links in the chain through which your company provides goods or services.
2. **A** To verify the security status of hardware and software components, you should begin running security tests on those components. Requesting security documentation on any hardware or software components received is useful but may not give you any added confidence in their security posture, since documentation can be forged or incomplete. Installing some hardware and software into critical systems is not the best choice since security scans of those systems may not identify compromised components. Additionally, contracting with a third-party assessor/monitor may be cost prohibitive.
3. **C** You should include minimum security requirements, as well as security responsibilities, in the new contract. If written correctly, these minimum security requirements will cover the other choices.
4. **A** In addition to data protection requirements imposed by regulations, you should also have the organization, and its personnel, sign nondisclosure agreements to ensure that sensitive data is protected and not disclosed to unauthorized parties. While data protection requirements may be included in the service level agreement, these may be more general and not enforceable on individuals that work for the third-party provider. Assessment reports may give you some insight into the provider's security posture but will not guarantee protection of sensitive data.



Objective 1.13 Establish and maintain a security awareness, education, and training program

In this objective we will discuss the organization's security awareness, education, and training program. This is one of the key administrative controls an organization has at its disposal, and the one that may be the most critical in protecting its assets.

Security Awareness, Education, and Training Program

An organization's security awareness, education, and training program is a key administrative control that addresses the most vulnerable element of the organization: the people. This program communicates to employees (the students) the threats, vulnerabilities, and risks associated with the human element. It's important to note that these terms—awareness, education, and training—are distinct, although sometimes used interchangeably. Table 1.13-1 details the unique elements of these three terms.

Methods and Techniques to Present Awareness and Training

Different levels of awareness, training, and education often require different presentation methods and can depend on factors such as the level of audience knowledge or understanding, the student's security responsibilities, depth of the subject, and so on. Frequency also may

TABLE 1.13-1 Training Levels and Descriptions

Term	Description	Examples
Awareness	Gives the student basic information on a topic; the "what" of the subject	Basic threats, vulnerabilities, and risk associated with the human element of an organization, such as social engineering
Education	Advanced or comprehensive instruction on a topic designed to give a student insight and understanding; the "why" of a subject	Advanced security knowledge or theory, such as encryption algorithm construction
Training	Gives the student intermediate knowledge of a topic and imparts skills; the "how" of the subject	Security-related skills, such as secure software configuration

dictate the type of presentation the student receives; for example, initial training and education is often more comprehensive than refresher awareness briefings or training, and thus may use different presentation methods. Logistical factors such as available equipment or training space may sometimes influence the presentation method used for the audience.

Presentation Techniques

Traditional presentation techniques, such as in-class training, may be preferred but not possible due to the size of the target audience, available space, remote offices, training budget, and so on. A popular form of training today is self-study, which could include prerecorded audio or video that the student can review on their own and recommended or required texts (e.g., books or websites). Another method of training that has increased in use over the past several years is distance learning using collaborative software over the Internet. This online learning offers the advantage of being able to reach a greater number of students, employs a live or synchronous training method, and may not be overly restricted by budget, training space, or distance.

Taking into account the best presentation methods to benefit learners means that some of the more traditional techniques, such as simply presenting a slide presentation, may not be effective for all students. Presentation techniques that are more interactive generally increase a student's retention of any information presented. Some of these interactive techniques include

- Social engineering role-play exercises
- Organizational phishing exercises
- Interactive security awareness training games with teams of students ("gamification")

In addition to presentation techniques, security topics should be tailored for the specific audience. Users with only very basic security responsibilities should be given a basic awareness overview when onboarding in the organization and at regular occurrences thereafter. Users that have more advanced security responsibilities, such as IT or security personnel, should be given more in-depth training, on more advanced topics, and at a more frequent rate. Even managers and senior executives should be presented with specific training that targets their unique roles and responsibilities, such as security risk, compliance, and other higher-level topics.

Often an employee may take the responsibility on of spearheading a training program or project and serve as the "security champion" for the project, leading others to adopt the security aspects of the project to integrate and improve security into their own areas. These security champions don't always have to be employees with security related duties; this shows that they have imbued the security concepts provided by extensive training and ensure that security becomes a built-in part of their worklife.



EXAM TIP Note that security topics can be presented in different ways, depending upon the level of knowledge or comprehension required, the audience, and the nature of the topic itself. Security topics can be presented simply as bulletin board notices, monthly newsletters, or in-depth classroom training. The presentation method should be adjusted to meet the needs of the organization.

Periodic Content Reviews

As threats, vulnerabilities, and risk continually change, the training presented to students in these areas should also occasionally change. When new threats and vulnerabilities become mainstream, employees need to be educated, perhaps in a focused training session, and the regular security training should be updated accordingly. Updates lend themselves to what is known as *just-in-time* training, meaning the student receives the training as soon as they need it. Additionally, training should be reviewed and updated when the organization has any type of operating environment change that could also amend its security environment, such as imposition of new governance or regulations on the organization, implementing new systems or security features, turnover in personnel, and so on.

Program Effectiveness Evaluation

The organization should be able to measure the effectiveness of its training program. Evaluation methods include collecting data following a training program on the increase or decrease in the number of security incidents or the number of suspected phishing incidents being reported. The key is to measure the specific training presented with expected results, to see if the training positively changes behavior, resulting in fewer security incidents and increased security compliance from employees.

The organization should record who gets trained, how often, and on what topics, and this documentation should be maintained in individual students' HR files. Ensuring users get the right type of training when they need it also serves to protect the organization if a user violates a policy or performs a malicious act. The person cannot later claim that they were not properly trained or even informed of the requirements.

REVIEW

Objective 1.13: Establish and maintain a security awareness, education, and training program In this objective we discussed security awareness, training, and education. Security awareness provides basic information on security topics, including threats, vulnerabilities, and risk, as well as basic security responsibilities an individual has in an organization.

Security training is normally targeted at developing skills, such as those that an IT or security person might require to perform their job functions. Security education presents topics that are advanced in nature and are geared toward higher-level understanding and comprehension of security subjects.

Presentation methods are critical, depending on factors such as the target audience, logistics (e.g., available personnel, space, and distance), and training budget. Traditional presentation methods such as classroom training can still be used, but other methods, based on the material presented and the knowledge level of the student, should be considered. These other methods include distance learning, self-study, and interactive simulations and exercises. Training should be evaluated and updated periodically for currency and relevancy to the organization. Just-in-time training should be considered for perishable knowledge or significant changes in threats and vulnerabilities.

The security awareness and training program should be evaluated periodically to determine its effectiveness; this is usually based on a measurement of how the training changes the behaviors of its target students. Results of an effective training program should lead to a decrease in security incidents and an increase in security-focused behaviors and compliance.

1.13 QUESTIONS

1. You have been asked by your supervisor to present a security topic to a small group of users in your company. All of the users work in the same building as you and there are plenty of conference rooms available for a short presentation. The topic is very basic and involves information regarding a new type of social engineering method used by attackers. Which level of instruction and presentation technique should you use?
   - **A.** Awareness, classroom training
   - **B.** Education, distance learning
   - **C.** Training, distance learning
   - **D.** Awareness, self-study
2. You are a cybersecurity analyst who works at a manufacturing company. Because of your experience and attendance at advanced firewall training, you are considered the local expert on the company's firewall appliances. Your supervisor has just told you that the company CISO wants you to give some training to some of the other cybersecurity technicians. Many of these technicians are geographically dispersed and on different work shifts. Which of the following would be the most effective way of presenting this training?
   - **A.** One-on-one training
   - **B.** Classroom training
   - **C.** Combination of distance-learning and self-study
   - **D.** Self-study

3. As part of governance requirements, a specific population of users within your company must be trained on risk management techniques and compliance on an annual basis. Which of the following groups is the likely target of this training?
   - **A.** Senior executives
   - **B.** Routine users
   - **C.** Help desk technicians
   - **D.** Junior security analysts
4. For a company-sponsored, off-duty employee development program, you would like to teach a more advanced class on security theory. Which of the following would be the appropriate level of this type of instruction?
   - **A.** Awareness
   - **B.** Training
   - **C.** Education
   - **D.** Briefing

1.13 ANSWERS

1. **A** Since all the users are co-located, distance learning may not be necessary. The topic is very basic so the training only needs to increase awareness about the new social engineering technique.
2. **C** Because some of the students you must train are geographically separated and work different shifts, classroom training likely won't be feasible. You should consider distance-learning, in combination with self-study, because of the advanced nature of the topic, and include interactive exercises to help the students learn how to configure the firewall better. Self-study alone likely would not enable them to learn these skills sufficiently.
3. **A** Since the topics involve risk management and compliance, senior executives likely benefit most from this type of training, as it is more suited to their roles and responsibilities.
4. **C** Advanced topics, such as security theory, would most likely be considered at the level of education within the training program, as this level of learning represents more advanced topics that cover the "why" of the subject.

Asset Security



Domain Objectives

- **2.1** Identify and classify information and assets.
- **2.2** Establish information and asset handling requirements.
- **2.3** Provision resources securely.
- **2.4** Manage data lifecycle.
- **2.5** Ensure appropriate asset retention (e.g., End-of-Life (EOL), End-of-Support (EOS)).
- **2.6** Determine data security controls and compliance requirements.

Domain 2, "Asset Security," discusses one of the most important things an organization can do in terms of securing itself: protect its assets. As you will see in this domain, anything of value can be categorized as an asset, and an organization must manage its assets effectively by identifying them and protecting them according to how valuable they are to the organization. In this domain we will discuss how to identify and classify assets and how to handle their secure use, storage, and transportation. We will also talk about how to provision assets and resources securely to authorized users. As an asset, information has a defined life cycle, and we will examine this life cycle so that you understand everything involved with protecting critical or sensitive information. We'll also discuss asset retention, as well as the security controls and compliance requirements that are levied as a part of protecting assets.



Objective 2.1 Identify and classify information and assets

As discussed in Objective 1.10, an *asset* is anything of value to the organization. Examples of assets include servers, workstations, network equipment, facilities, information, and people. The value placed on an asset depends on how much the organization needs it to fulfill its mission. How an organization places value on an asset depends on many factors, including the original cost of the asset, its depreciation, its replacement value, what it costs to repair, how much revenue it generates for the organization, and how critical it is to the organization's mission.

We also discussed tangible and intangible assets in Objective 1.10. As a refresher, a tangible asset is an item that we can easily see, touch, interact with, and measure; examples are systems, equipment, facilities, people, and even information. Intangible assets are those that are difficult to interact with, and sometimes difficult to place a value on, but are still critical to the organization's success. Examples of intangible assets include customer satisfaction, reputation, and standing in the marketplace. In this objective we will discuss how assets and data are identified and classified.



EXAM TIP Although the CISSP objectives sometimes seem to refer to *information* and *assets* as two different things, understand that in reality information is in fact one of the organization's critical assets, and in this book, we will treat it as such and make reference to it as an asset, along with facilities, equipment, systems, and even people.

Asset Classification

Classifying assets means to categorize them in some fashion so that they can be managed better. Assets are generally classified in terms of criticality and sensitivity, both of which are discussed in the next section. Assets aren't purchased or acquired for their own sake; they exist

to support the mission of the organization, including its business processes. The organization can classify an asset according to criticality by determining how critical the business process is that the asset supports.

A common way to perform this criticality assessment is to perform a *business impact analysis (BIA)*, introduced in Objective 1.8. A BIA is also one of the first steps in business continuity planning, which we will discuss later in the book. A BIA inventories the organization's business processes to determine those that are critical to its mission and cannot be allowed to stop functioning. It also identifies the information assets used to support those processes, which are considered the critical assets in an organization.

Cross-Reference

We will discuss business impact analysis, as part of business continuity planning, in Objective 7.13.

Even before determining criticality, however, the organization must discover all of its assets. Normally, tangible IT/cyberassets fall into one of two categories: hardware or software. Hardware is normally inventoried and identified by its serial number, model number, cost, and other data elements. Software is tracked by its license number or key, business or process it supports, and specific function.

Data Classification

Like other assets, data can be classified in terms of criticality, or in terms of sensitivity. While often used interchangeably, data criticality and data sensitivity mean two different things. Data criticality refers to how important the data is to a mission or business process, and data sensitivity relates to the level of confidentiality of the data. For example, privacy data in an organization may not be critical to its core business processes and only collected incidentally to the human resource process of managing employees. However, it is still very sensitive data that cannot be allowed to be accessed by unauthorized entities.

Cross-Reference

Remember from Objective 1.2 that *data* and *information* are terms that are often used interchangeably, but they have distinct definitions. Data are raw, singular pieces of fact or knowledge that have no immediate context or meaning. An example would be an IP address, or domain name, or even an audit log entry. Information is data organized into context and given meaning. When given context, many pieces of data become information. Giving data context means correlating data—determining how it relates to other pieces of data and a particular event or circumstance. While individual pieces of data should be considered critical or sensitive, we normally classify information rather than pieces of data.

Also like other assets, information must be identified and inventoried. Although information is a tangible asset, sometimes it may seem more like an intangible asset because inventorying it is much more challenging than simply counting computers in an office. The organization has to determine what information it has before it can figure out its criticality and sensitivity.

Identifying information usually involves looking at every system in the organization and recording what types of information are processed by them. This could be different sets of information that span privacy, medical, and financial information, or information that relates to proprietary processes the organization uses to stay competitive in the marketplace.

Once the different types of information are identified and inventoried, those different types are inventoried again by where they are located, on which system(s) they are processed, and how the systems process that data type. They can then be assigned criticality and sensitivity values. Remember from our discussion in Objective 1.10 that we can assign quantitative or qualitative values to different elements of risk. We can also assign quantitative or qualitative values to information. Quantitative values are normally numerical, such as assigning a dollar amount to the information type. Quantitative values in dollars are easy to assign to a piece of equipment, but not necessarily to the type of information it processes. Qualitative values may be better suited for information and can range, for example, from very low to very high.

The organization should develop a formalized classification system for its assets. Some assets may be assigned a dollar value for criticality to the organization, but other assets may have to be assigned a classification range that is qualitative in nature. Unlike quantitative values, qualitative values can be assigned for both criticality and sensitivity. This classification system should be documented in the organization's data classification and asset management policies.

One of the good reasons for assigning assets to a classification system is to determine the cost effectiveness of assigning controls to protect those assets. We discussed the criteria for selection of controls and countermeasures in Objective 1.10. The value (in terms of cost, criticality, or sensitivity) of an asset, whether it is an individual server or an information type, must be balanced with the cost of a control that is implemented to protect it. If the protection is insufficient or the asset's value is far less than the cost to implement and maintain the control, then implementing the control is not worthwhile. Implementing the control should also be balanced with the level of risk the control mitigates or reduces. If an asset's classification is very high for criticality and/or sensitivity, then obviously it is more valuable to the organization and the cost/benefit analysis will support implementing a more costly control to protect it.

Various different classification systems are used across private industry and government. Private companies may use a classification system that assigns labels such as Public, Private, Proprietary, and Confidential to denote asset sensitivity, or subjective values of Very Low, Low, Moderate, High, and Very High to denote either sensitivity or criticality. Military classifications typically include the use of the terms Top Secret, Secret, and Confidential to classify information and other assets according to sensitivity. Regardless of the classification system used, the organization must formally develop and document the system, and apply it when determining the level of protection that assets require.

Cross-Reference

We will discuss classification labels in more detail in Objective 2.2 when we cover data handling requirements.

REVIEW

Objective 2.1: Identify and classify information and assets In this objective we reviewed the definitions of tangible and intangible assets. We also discussed the need for identifying and classifying assets, including information. Assets can be classified according to criticality and sensitivity. Criticality describes the importance of the asset to the mission or business, and may be determined by examining critical business processes and further determining which assets support them. Sensitivity relates to the need to keep information or assets confidential and away from unauthorized entities. Assets must be identified and inventoried before they can be classified according to criticality and sensitivity.

2.1 QUESTIONS

1. You are tasked with inventorying assets in your company. You must identify assets according to how important they are to the business processes and the company. Which of the following methods could you use to determine criticality?
   - **A.** Classification system
   - **B.** Business impact analysis
   - **C.** Data context
   - **D.** Asset value
2. Which of the following is the best example of information (data with context)?
   - **A.** IP address
   - **B.** TCP port number
   - **C.** Analysis of connection between two hosts and the traffic exchanged between them
   - **D.** Audit log security event entry
3. You are creating a data classification system for the assets in your company, and you are looking at subjective levels of criticality and sensitivity. Which of the following is an example of a subjective scale that can be used for data sensitivity and criticality?
   - **A.** Very Low, Low, Moderate, High, and Very High
   - **B.** Dollar value of the asset
   - **C.** Public, Private, Proprietary, and Confidential
   - **D.** Top Secret, Secret, and Confidential
4. Your company is going through the process of classifying its assets. It is starting the process from scratch, so which of the following best describes the order of steps needed for asset classification?
   - **A.** Classify assets, identify assets, and inventory assets
   - **B.** Identify assets, inventory assets, and classify assets
   - **C.** Classify assets, inventory assets, assign security controls to assets
   - **D.** Identify assets, assign security controls to assets, classify assets

2.1 ANSWERS

1. **B** Conducting a business impact analysis (BIA) can help you to not only identify critical business processes but also identify the assets that support them and assist you in determining asset criticality.
2. **C** Data are individual elements of knowledge or fact, such as an IP address, TCP port number, or audit log security event entry without context. An example of information (data placed into context) would be an analysis of a connection event between two hosts and the traffic that is exchanged between them, which would consist of multiple pieces of data.
3. **A** Subjective values are typically nonquantitative and offer a range that can be applied to both criticality and sensitivity, in this case a qualitative scale from very low to very high. Dollar value of the asset would be a quantitative or objective measurement of criticality and would not necessarily apply to sensitivity of assets. Commercial sensitivity labels, such as Public, Private, Proprietary, and Confidential, would not necessarily apply to criticality of assets. Top Secret, Secret, and Confidential are elements of government or military classification systems and do not necessarily denote criticality, nor would they apply to a commercial company.
4. **B** The correct order of steps is identify the assets, inventory the assets, and classify the assets by sensitivity and criticality. Assigning security controls to protect an asset comes after its criticality and sensitivity are classified.



Objective 2.2 Establish information and asset handling requirements

This objective is a continuation of our discussion on how to classify and handle sensitive information assets (both information and the systems that process it). In Objective 2.1 we discussed how assets in general, and information specifically, are categorized (or classified) by criticality and sensitivity. This objective takes it a little bit further by discussing how those classification methods affect how an asset is handled during its life.

Information and Asset Handling

An organization's data sensitivity or classification policy assigns levels of criticality or sensitivity to different types of information (e.g., private information, proprietary information, and so on). Controls or countermeasures are assigned to protect sensitive and critical data based on those classification levels. Assets with higher criticality or sensitivity classification labels

obviously get more protection because of their greater value to the organization. We're going to further that discussion in this objective by focusing on how information is actually handled as a result of those classifications.



EXAM TIP Remember that assets are classified according to criticality and/or sensitivity.

Handling Requirements

Information has a life cycle, which we will discuss at various points in this book. For now, the stages of the life cycle that you should be aware of are those that involve how information is handled during storage (at rest), physical transport, transmission (in transit), and when it is transferred to another party. We will discuss other aspects of the information life cycle later, to include generating information as well as its disposal.

The following general handling requirements apply to information assets regardless of where they are in their life cycle:

- Encryption for information in storage and during transmission
- Proper access controls, such as permissions
- Strong authentication required to access information assets
- Physical security of information assets
- Positive control and inventory of assets
- Administrative controls, such as policies and procedures, which dictate how information assets are handled

We will now examine how these handling requirements apply to information assets in various stages, such as storage, transportation, transmission, and transfer.

Storage

Information in storage, often referred to as *data at rest*, is information that is not currently being accessed or used. Storage can mean that it resides on a hard drive in a laptop waiting to be processed or used, or stored in a more permanent environment for archival purposes or as part of a backup/restore strategy. It's important to make the distinction between information that has been archived and information that is simply backed up for restoration in the event of a disaster.

Archived information is often what the organization may be required to keep for a specified time period due to regulatory or business requirements, but that information will not be used to restore data in the event of an incident or disaster. The organization's information retention policies, articulated with external governance requirements, typically indicate how long

information must be archived. In general, it is a good practice to archive information only to the minimum extent mandated by regulation or business requirements. Keeping information archived after it is no longer required to be retained is not a good practice because it may eventually be compromised or misused. Information that is backed up to restore data that is lost is stored on backup media locally on site or, preferably, off site, such as in the cloud or an offsite physical storage location.

Regardless of whether information is stored for archival or backup purposes, specific protections should be in place for handling:

- Encryption of backup or archival media
- Access permissions to backup or archival media
- Strong authentication required to access media
- Physical security of archived or backed up information, to include locked or secure storage, physical access logs, proper markings, and so on
- Inventory of information stored as well as the media and devices that store it
- Administrative controls, such as archival and backup/restore procedures, two-person control, and so on

Transportation

Physically transporting information assets, such as archival or backup media, decommissioned workstations, or new assets that are transferred to another geographic location within the organization, requires special handling. In addition to the controls we've already mentioned, such as encrypting media and requiring strong authentication to access media and devices, the physical transportation of information assets may include additional physical controls. Maintaining positive control by any individuals assigned to transport sensitive information assets is one such consideration; information assets may require constant presence of someone assigned to escort or guard the asset while it is being transported, as well as a strong chain of custody during transport.

Transmission

Information that is digitally transferred, also known as *data in transit*, is typically protected by the use of strong authentication controls and encryption. These measures ensure that information is not intercepted during transmission, and that it is sent, received, and decrypted only by properly authenticated individuals.

Cross-Reference

We'll discuss transmission security in more depth in Domain 3.

Transfer

When information assets are transferred between individuals or parties, it usually involves a physical transfer of media or systems. However, logical transfer between entities, such as transmitting protected health data, also involves specific protections unique to transferring information assets. All of these protection methods have already been discussed but may apply uniquely to instances of transferring assets, such as:

- Encrypting storage or transmission media prior to transfer
- Implementing strong authentication controls to ensure that personnel are authorized to send or receive information assets
- Positive control and inventory of physical assets, such as systems
- Administrative, technical, and physical procedures for verifying identities of transferring parties

Information Classification and Handling Systems

As you can see from the previous discussion, information classification schemes directly dictate how information assets are handled, both physically and logically. Information classification and handling schemes are usually categorized in two ways for the CISSP exam: those that apply to commercial or private organizations and those that are used by government and military organizations. We will discuss both briefly, as well as offer examples of each.

Commercial/Private Classification Handling

Commercial and private organizations have a myriad of information classification and handling systems, so the classification levels we discuss here only serve as examples. Because of the wide variety of information types, organizations may have many of these classification levels, which apply to handling sensitive (but not necessarily critical) information assets. Table 2.2-1 lists some of these classification levels as examples.

TABLE 2.2-1 Examples of Commercial and Private Information Classification and Handling Labels

Classification	Description
Public	Nonsensitive information that is considered releasable to the general public under specific circumstances
Private	Information that may be considered privacy information, such as an individual's health, financial, or personal information
Confidential or sensitive	Information that may be considered sensitive by the organization, such as financial information, personnel actions, and so on
Proprietary	Information that may be key to maintaining a competitive place in the market, such as formulas, processes, and so on

TABLE 2.2-2 U.S. Government Classification and Handling Levels

Classification	Description
Top Secret	Information assets whose unauthorized disclosure could result in exceptionally grave danger to the nation.
Secret	Information assets whose unauthorized disclosure could result in serious damage to the nation.
Confidential	Information assets whose unauthorized disclosure could cause damage to the nation.
Unclassified	While not a true classification level, this information category is a catch-all and applies to any information that is not considered "classified" (i.e., confidential, secret, or top secret), and includes restricted information, privacy or financial information, law enforcement information, etc. Unauthorized disclosure of these information assets could cause undesirable effects if available to the public.

U.S. Government Classification and Handling

Many government agencies have their own classification and handling systems, and one of the most well-known is the system used by the U.S. government and the U.S. armed forces. Table 2.2-2 describes those classification and handling levels. Note that this system describes the relative sensitivity of information and systems, not criticality.

Note that different information types may have handling requirements prescribed based on their sensitivity levels and regulations that address those information types. For example, protected health information (PHI), which is addressed by the Health Insurance Portability and Accountability Act (HIPAA), has its own special handling requirements.

REVIEW

Objective 2.2: Establish information and asset handling requirements In this objective we looked at the different aspects of information asset handling, including storage, transportation, transmission, and transfer. We examined different methods of protecting information as it is being handled, such as encryption, strong authentication, access controls, physical security, and other administrative controls. We also discussed different information handling requirements, such as those that a private company may have and those of government agencies and military services. These include the different handling requirements for information designated as Public, Private, Confidential, and Proprietary (labels often used in the private or commercial sector) and information designated as Top Secret, Secret, Confidential, and Unclassified (used, for example, in the U.S. government and military services).

2.2 QUESTIONS

1. You are a cybersecurity analyst in your company and you have been tasked with ensuring that hard drives containing highly sensitive information are transported securely to one of the company's branch offices located across the country for use in its systems there. Which of the following must be implemented as a special handling requirement to ensure that no unauthorized access to the data on the storage media occurs?
   - **A.** Handling policies and procedures
   - **B.** Storage media encryption
   - **C.** Data permissions
   - **D.** Transmission media encryption
2. You are a cybersecurity analyst working for a U.S. defense contractor, so you must use government classification schemes to protect information. Which of the following classification labels must be used for information that may cause serious damage to the nation if it is disclosed to unauthorized entities?
   - **A.** Top Secret
   - **B.** Proprietary
   - **C.** Secret
   - **D.** Confidential

2.2 ANSWERS

1. **B** While each of these is important, encrypting the storage media on which the information resides is critical to ensuring that no unauthorized entity accesses the data during the physical transportation process. Information handling policies and procedures should include this requirement. Data permissions are only important after the media reaches its destination since authorized personnel will then be able to decrypt the drive and, with the correct permissions, access specific data on the decrypted media. Since the media is being physically transported, transmission media encryption is not an issue here.
2. **C** Information whose unauthorized disclosure could cause serious damage to the nation is labeled as Secret in the classification scheme used by the U.S. government and military services.



Objective 2.3 Provision resources securely

In this objective we discuss how to ensure that organizational information resources are made available to the right people, at the right time, and in a secure manner. This is only made possible through strong security of information assets and good asset management, both of which we will discuss in this objective.

Securing Resources

Asset security is about protecting all the resources an organization deems valuable to its mission. This is true whether the assets are data, information, hardware, software, equipment, facilities, or even people. Information assets (which consist of both information and systems) must be secured by a variety of administrative, technical, and physical controls. Examples of these controls include encryption, strong authentication systems, secure processing areas, and information handling policies and procedures. In order to secure these resources, we must consider items such as information asset ownership, the control and inventory of assets, and how these assets are managed.

Asset Ownership

Asset ownership refers to who is responsible for the asset, in terms of maintaining it and protecting it. Remember that information assets are there to support the organization's business processes. The asset owner may or may not be the owner of the business process that the asset supports. In any event, the asset owner must ensure that the asset is maintained properly and that security controls assigned to protect the asset are properly implemented.

Assets can be assigned owners based on the critical business processes the assets support or other functional or organizational structures, such as the function of the business process or the department that funds the asset. In any event, asset owners must be aware of their responsibilities in maintaining and securing assets.



EXAM TIP Asset and business process owners may not be the same, as assets may be owned by another entity, especially when an asset supports multiple business processes.

Asset Inventory

As mentioned in Objective 2.1, an asset inventory can be developed by performing a business impact analysis, which identifies and inventories critical business processes and the assets that support those processes.

In their simplest form, inventories of equipment, particularly hardware and facilities, are relatively easy to maintain but must be consistently managed. Some aspects of maintaining a good asset inventory include

- Tracking assets based on asset value, criticality, and sensitivity
- Assigning to assets owners who are responsible for their maintenance and security
- Maintaining knowledge of an asset's location and status at all times
- Giving special attention to high-value assets, including tracking by serial number, protection against theft or misuse, and more frequent inventory updates

There may be some differences in how the organization assigns value to, inventories, and tracks tangible assets versus intangible assets, as discussed next.

Tangible Assets

Tangible assets, as previously described, are those that an organization can easily place a value on and interact with directly. These include information, pieces of hardware, equipment, facilities, software, and so on. These assets can more easily be inventoried and should be tracked from multiple perspectives, including the following:

- Securing the asset from theft, damage, or misuse
- Ensuring critical maintenance activities are performed to keep the asset functional
- Replacing component parts when needed
- Managing the asset's life cycle, which includes acquisition, implementation, and disposal phases

Intangible Assets

Managing intangible assets is not as easy as managing tangible ones. Remember that intangible assets are those that cannot be easily assigned a monetary value, counted in inventory, or interacted with. Examples of intangible assets include things that are nebulous in nature and may radically change in value or substance at a moment's notice, such as consumer confidence, standing in the marketplace, company reputation, and so on. While there are many qualitative measurements that can be conducted to determine the relative point-in-time value of these intangible assets, such as statistical analysis, sampling, surveys, and so on, these measurements are simply good guesses at best. Additionally, since the value of intangible assets fluctuates frequently and sometimes wildly, these measurements must be constantly reassessed. An organization can't simply "inventory" consumer confidence once and expect that it will be the same the next time it is measured; these assets must be measured constantly.

Asset Management

Asset management is not only about maintaining a good inventory. Overall asset management ensures that resources, including information assets, are provided to the people and processes that need them, when they need them, and in the right state, in terms of function

and performance. A critical part of asset management is secure provisioning, which involves providing access to information assets only to authorized entities. *Secure provisioning* is the collection of processes that ensure that only authorized, properly identified, and authenticated entities (usually individuals, processes, and devices) access information and systems and can only perform the actions granted to them. Secure provisioning typically involves

- Verifying an entity's identity (authentication)
- Verifying the entity's need for access to systems and information (authorization)
- Creating the proper access tokens (e.g., usernames, passwords, and other authentication information)
- Securely providing entities with authentication and access tokens
- Assigning the correct authorizations (rights, permissions, and other privileges) for entities to access the resource based on need-to-know, job position or role, security clearance, and so on
- Conducting ongoing privilege management for assets

Note that secure provisioning is only one small part of end-to-end asset management. As mentioned, we will discuss the full asset management life cycle in later objectives.

REVIEW

Objective 2.3: Provision resources securely In this objective we further discussed managing information assets. We examined the concepts of information asset ownership and noted that the business process owner the asset supports is not always the same individual who owns the asset. Asset owners have the responsibility of maintaining and securing any assets assigned to them. Assets must be inventoried and tracked based on criticality, sensitivity, and value to the organization, whether that value is expressed in monetary terms or not. Tangible assets are much easier to manage because they can be physically counted, tracked, located, and assigned a dollar value. Intangible assets are more difficult to manage and can change in value frequently. Intangible assets include consumer confidence, company reputation, and marketplace share. Intangible assets can be measured using qualitative or subjective methods but must be measured constantly.

Asset management involves not only inventory of assets, but also providing assets to the right people, at the right time, and in the state they are required. This starts with secure provisioning of assets. Secure provisioning includes creating access tokens and identifiers, enforcing the proper identification and authentication of entities, and authorizing actions an entity can take with an asset.

2.3 QUESTIONS

1. Your supervisor wants you to assign asset values to certain intangible assets, such as consumer confidence in the security posture of your company. You consider factors such as risk levels, statistical and historical data relating to security incidents within the company, and other factors. You also consider factors such as cost of various aspects of security, including systems and personnel. Which of the following is the most likely type of value you could assign to these assets?
   - **A.** Quantitative values based on customer confidence surveys
   - **B.** Quantitative values related to cost versus revenue
   - **C.** Qualitative descriptive values of very low, low, moderate, high, and very high
   - **D.** Qualitative values related to cost versus revenue
2. You are a cybersecurity analyst who is in charge of routine security configuration and patch management for several critical systems in the company. One of the systems you work with is a database management system that supports many different departments and lines of business applications within the company. Which of the following would be the most appropriate asset owner for the system?
   - **A.** Information technology manager for the company
   - **B.** Asset owner who is chiefly responsible for maintenance and security of the asset
   - **C.** Multiple business process owners sharing asset ownership
   - **D.** Business process owner for the most critical line of business application

2.3 ANSWERS

1. **C** Any value assigned to intangible assets will be qualitative in nature; subjective descriptive values, such as very low, low, moderate, high, and very high are the most appropriate values for intangible assets since it is difficult to place a monetary value on them.
2. **B** Since the asset supports many different business processes and owners, the system should be owned by someone who is directly responsible for maintenance and security of the asset; however, they may be further responsible and accountable to any or all of the business process owners.



Objective 2.4 Manage data lifecycle

In this objective we will discuss the phases of the data life cycle, as well as the roles different entities have with regard to using, maintaining, and securing data throughout its life cycle. We will also discuss different interactions that can occur with information, such as its

collection, maintenance, retention, and destruction. Additionally, we will look at aspects of data that should be closely monitored, including its location and data remanence.

Managing the Data Life Cycle

Data has a formalized life cycle, which starts with its creation or acquisition. There are several data life cycle models that exist, but a commonly referenced one from the U.S. National Institute of Standards and Technology (NIST). NIST Special Publication (SP) 800-37, Revision 2, defines the information life cycle as "The stages through which information passes, typically characterized as creation or collection, processing, dissemination, use, storage, and disposition, to include destruction and deletion." The rest of this objective discusses the roles that individuals and other entities play during the data life cycle and how they are involved in various data life cycle activities.

Data Roles

Managing an organization's data requires assigning sensitivity and criticality labels to the data, assigning the appropriate security measures based on those labels, and designating the organizational roles responsible for managing data within the organization. While an organization may not formally assign data management roles, some regulations require formal data ownership roles (e.g., the General Data Protection Regulation, or GDPR, as covered later in this objective). Data management roles can be held legally accountable and responsible for the protection of data and its disposition. Note that many of these roles can overlap in role and responsibility unless governance specifically prohibits it.



ADDITIONAL RESOURCES NIST SP 800-18, Revision 1, details different data management roles.

Most data management roles must be formally appointed by the organization, and access to data may be granted only if an entity meets certain criteria, based on regulations, including the following:

- Security clearance
- Need-to-know
- Training
- Statutory requirements

Data Owners

A *data owner* is an individual who has ultimate organizational responsibility for one or multiple data types. A data owner typically is an organizational executive, such as a senior manager,

a vice president, department head, and so on. It could also be someone in the organization who has the legal accountability and responsibility for such data.

The data owner and asset owner may not necessarily be the same individual (although they could be, depending on the circumstances); remember from Objective 2.3 that an asset owner is the individual who is accountable and responsible for assets or systems that process sensitive data. Both data and system (asset) owners typically have the following responsibilities:

- Develop a system or data security plan
- Identify access requirements for the asset
- Implement and manage the security controls for systems or data

Controllers, Custodians, Processors, Users, and Subjects

While there are many commonsense roles associated with managing data in an organization, there are some other key roles you should understand. Some roles are defined and required by regulation, such as a *data controller*. A data controller is the designated person or legal entity that determines the purposes and use of data within an organization. This is normally a senior manager in most organizations.

A *data processor* is a person who routinely processes data in the performance of their duties; this may be someone who performs simple data entry tasks, views data, runs queries on data, and so forth. Note that this is a generalized description of the data processor, but one regulation in particular, the European Union's General Data Protection Regulation (GDPR), specifically defines a data processor as someone who provides services to and processes data on behalf of the data controller. Data processors don't have to be individuals; a third-party service provider contracted to process data on behalf of a data controller also fits this definition.



NOTE Another role that GDPR requires is the Data Protection Officer, which is a formal leadership role responsible for the overall data protection approach, strategy, and implementation within an organization. This person is also normally accountable for GDPR compliance.

A *data custodian* is a more generalized role and is an appointed person or entity that has legal or regulatory responsibility for data within an organization. An example would be appointing someone in a healthcare organization as the custodian for healthcare data. This might be a data privacy officer or other senior role; however, all data management roles have some custodial responsibility to varying degrees.

Data users are simply those people or entities that use data in their everyday job. They could simply access data in a database, perform research, or conduct data entry or retrieval.



EXAM TIP Remember that a data user is someone who uses the data in the normal course of performing their job duties; a data subject is the person about whom the data is used.

Finally, a *data subject* is the person whose personal information is collected, processed, and/or stored. If a person submits their personal, healthcare, or financial information to another entity, that person is the subject of the data. The overall end goal of privacy is to protect the subject of the data from the misuse or unauthorized disclosure of their data.



NOTE The EU's GDPR formally defines critical roles in managing data, including the controller, processor, and subjects.

Data Collection

Data collection is an ongoing process which occurs during routine transactions, such as inputting medical data, filling out an online web form with financial data, and so on. Data collection can be formal or informal and performed via paper or electronic methods. The key to data collection is that organizations should only collect data they need either to fill a specific business purpose or to comply with regulations, and no more. This constraint should be expressed in the organization's data management policies (e.g., data sensitivity, privacy, etc.).

Data Location

Data location refers to the physical location where the data is collected, stored, processed, and even transmitted to or from. Data location is an important consideration, since data often crosses state and international boundaries, where laws and regulations governing it may be different. Some countries exert a concept known as *data sovereignty*, which essentially states that if the data originates from or is stored or processed in that country, their government has some degree of regulatory oversight over the data.

Cross-Reference

Objective 1.5 discussed *data localization* laws (also called *data sovereignty* laws), which touch on the legal ramifications of data location.

Data Maintenance

Data maintenance is a process that should be performed on a regular basis. Data maintenance includes many different processes and actions, including

- Validating data for proper use, relevancy, and accuracy
- Archiving, backing up, or restoring data
- Properly transferring or releasing data
- Disposing of data

Data maintenance actions should be documented and only performed by authorized personnel, in accordance with the organization's policies and governing regulations.

Data Retention

Data retention comes in two forms:

- Backing up data so it can be restored in the event it is lost due to an incident or natural disaster. This is a routine business process that should be performed on critical data so that it can be restored if something happens to the original source of the data.
- Archiving data that is no longer needed. This involves moving data that is no longer used but is required for retention due to policy or legal requirements to a separate storage environment.

As a general rule, an organization should only retain data that it absolutely needs to fulfill business processes or comply with legal requirements, and no more; the more data that is retained, the higher the chances are of a breach or legal liability of some sort. The organization should create a data retention policy that complies with any required laws or regulations and details how long data must be kept, how it must be stored, and what security protections must be afforded it. Any retained data should be inventoried and closely monitored for any signs of unauthorized access or use.



EXAM TIP Understand the difference between a data archive and a data backup. A data archive is data that is no longer needed but is retained in accordance with policy or regulation. A data backup is created to restore data that is still needed, in the event it is lost or destroyed.

Data Remanence

Data remanence is that which remains on a storage media after the media has been sanitized of data. It could be simply random ones and zeros, file fragments, or even human- or machine-readable information if the sanitization process is not very thorough. Data remanence is a problem when data remains on media that is transferred to another party or reused. The best way to solve the data remanence problem is destruction of the media on which it resides.

Data Destruction

Data destruction refers not only to the process of destroying data itself by wiping data from storage media, but also destroying paper copies and any other media on which data resides. Media includes hard drives, USB sticks, optical discs, and so on. In the case of routine or

noncritical data destruction, it may simply be enough to degauss hard drives or smash SD cards and optical discs. If the data is very sensitive, often the destruction must be witnessed and documented by more than one person, with the process recorded and verified. The following are common methods of destroying sensitive data and the media on which it resides:

- Burning or melting
- Shredding
- Physical destruction by smashing media with hammers, crowbars, etc.
- Degaussing hard drives or other magnetic media

REVIEW

Objective 2.4: Manage data lifecycle In this objective you learned about the data life cycle, as well as some of the activities that go on during that life cycle. We also discussed the different roles that entities play in the data life cycle, such as data owners, controllers, custodians, processors, users, and data subjects. We also covered various activities that are critical to managing data during its life cycle, including its collection, storage or processing, location, maintenance activities that must be performed on data, how data is retained and destroyed, and data remanence.

2.4 QUESTIONS

1. You are drafting a data management policy in order to adhere to laws and regulations that apply to the various types of personal data your organization collects and processes. Which of the following would be the most important consideration in developing this policy?
   - **A.** Cost of security controls used to protect data
   - **B.** Appointment of formal data management roles
   - **C.** Manual or electronic methods of collecting data from subjects
   - **D.** Background check requirements for data users in the organization
2. Your organization has developed formalized processes for data destruction, some of which require witnesses for the destruction due to the sensitivity of the data involved. Which of the following is the organization trying to reduce or eliminate?
   - **A.** Data collection
   - **B.** Data retention
   - **C.** Data remanence
   - **D.** Data location

2.4 ANSWERS

1. **B** While all of these may be important considerations for the organization and its data management strategy, to be compliant with certain data protection regulations the organization must appoint formal data management roles. Cost of security controls isn't considered or addressed by policy, nor is the method of collecting data from subjects. Background check requirements for employees who use data in the organization are typically addressed by other policies.
2. **C** The organization, through a comprehensive destruction process, is attempting to reduce or eliminate any data remanence that may be compromised or accessed by unauthorized entities. Data collection is part of its business processes that it must also address, as is retention and location, but these are not addressed through data destruction.



Objective 2.5 Ensure appropriate asset retention (e.g., End-of-Life (EOL), End-of-Support (EOS))

In Objective 2.4 we discussed data retention; in this objective we will expand that discussion to examine asset retention in general. There are issues that an organization must consider during its asset life cycle, including an asset's normal end-of-life point and how to replace it and dispose of it.

Asset Retention

In our coverage of data retention in Objective 2.4 we discussed how the organization should only retain data that is necessary to fulfill business processes or legal requirements, and no more. Retaining unnecessary data presents a greater chance of compromise or unauthorized access. Similarly, an organization needs to consider retention issues related to its other assets, such as systems, hardware, equipment, and software, and develop a life cycle for those assets.



CAUTION Data should only be retained when required due to governance (i.e., laws, regulations, and policies) or to fulfill business process requirements. Retaining data unnecessarily subjects the organization to risk of data compromise and legal liability.

Asset Life Cycle

The asset life cycle is similar to the general data life cycle. Assets are acquired (software may be developed or purchased and hardware may be bought or constructed) and go through a life cycle that includes implementation, maintenance (sometimes called sustainment), and finally disposal. The asset life cycle typically contains the following phases (although there are many versions of the asset life cycle promulgated in different industries and communities):

- **Requirements** Establish the requirements for what the asset needs to do (function) and how well it needs to do it (performance)
- **Design and architecture** Establish the asset design and overall fit into the architecture of the organization
- **Development or acquisition** Develop or purchase the asset
- **Testing** Test the asset's suitability for its intended purpose and how well it integrates with the existing infrastructure
- **Implementation** Put the asset into service
- **Maintenance/sustainment** Maintain the asset throughout its life cycle by performing normal activities such as repair, patching, upgrades, etc.
- **Disposal/retirement/replacement** Determine that the asset has reached the end of its usable life and replace it with another asset, as well as dispose of it properly

As mentioned, these are generic asset life cycle phases; different asset management and systems/software engineering methodologies have similar phases but may be referred to differently.

End-of-Life and End-of-Support

Assets normally have a prescribed life cycle, at the end of which they are often replaced or must be disposed of. Additionally, some assets may outlive their useful life cycle because organizations are reluctant to replace assets that are costly or have a specialized purpose. Often, replacing an asset involves retooling or re-architecting an entire infrastructure. New assets, especially complex systems, often require additional training on their use and maintenance.

In any event, some assets are also no longer supported by the manufacturer or developer. This may be because the asset has outlived its useful life and there is a newer asset that can perform its function far better, more efficiently, or more securely. An example is the Windows 7 operating system, which lost its official support from Microsoft in January 2020. End-of-support means Microsoft will no longer provide patches, updates, or support for issues that result from still running that particular operating system version. However, even in 2022, there are still legacy systems running Windows 7. This is a good example of the difference between the terms end-of-life and end-of-support. End-of-life means that the asset no longer fulfills its intended purpose, is obsolete, broken, or the organization's need for it has passed. It must be replaced by a newer asset. End-of support, however, means that

the asset may still be functioning as intended, it still serves its purpose, and the organization will still use it, albeit at a greater risk, however, since its vendor no longer offers maintenance, upgrades, and so on.



EXAM TIP End-of-life means that the item is no longer viable or functional as required. End-of-support simply means that its manufacturer no longer services it, regardless of its level of functionality.

When an asset, such as software or a piece of hardware, exceeds its end-of-life point or its end-of-support time, the organization must decide how to best handle replacing the asset, if its function is still needed, and then what to do with the asset. Some assets can be donated to other organizations, and some can be repurposed, but many assets must be destroyed.

Replacement, Destruction, Disposal

Replacing an asset with one that performs similar functions is not the only consideration when the asset reaches its end-of-life. In addition to being able to perform similar functions as the asset it replaces, the organization should consider performance factors when replacing an asset. The asset not only must perform the original function but, ideally, should perform it more efficiently, effectively, and securely. An organization may require new features when it replaces an asset or it may require the asset to have certain characteristics that make it interoperable with newer technologies

Once an asset is replaced, the organization has to decide how to dispose of it. Often, donating an asset, such as a general-purpose computer system, may be an option, particularly if another organization can use it for purposes that don't require as much functionality or performance capability of the asset. (Think of an organization that replaces all of its desktops, and then donates those computers to a school or charity.) However, before the asset is repurposed or transferred to another organization, it must be cleared of all sensitive data; in fact, any data on the system should be purged so that it can never be accessed again by unauthorized entities, as discussed in Objective 2.4. This is where steps to remove all data remanence are performed; storage media such as hard drives are removed from the asset or, at a minimum, wiped to destroy data on the media.

Some assets can't be transferred or repurposed. Assets that are typically specialized use items may have applications that the organization does not wish to allow someone else to make use of. Consider an outdated but still sophisticated system that performs a particular manufacturing process that the organization does not want to be repeated by a competitor. Or think of systems with military-specific applications that should not be transferred to the general public. These types of systems will likely be destroyed rather than repurposed. They will still go through the typical steps of being purged of sensitive data, but then the hardware will generally be destroyed so that it can no longer be used or reverse engineered.

REVIEW

Objective 2.5: Ensure appropriate asset retention (e.g., End-of-Life (EOL), End-of-Support (EOS)) In this objective you learned about general asset retention principles, including the asset life cycle and how an asset's end-of-life and end-of-support points affect an organization's decisions regarding whether to replace the asset, repurpose it, or otherwise destroy it.

2.5 QUESTIONS

1. Your company has requirements for a new piece of engineering software that must have specific characteristics. The software is replacing an older software package that was developed internally and, as such, lacks many new features and modern security mechanisms. The organization has determined that it is not cost-effective to employ teams of developers to create a new application. In which phase of the asset life cycle would the organization be when replacing the software package?
   - **A.** Acquisition
   - **B.** Disposal
   - **C.** Implementation
   - **D.** Maintenance
2. You work for a company that is replacing all of its laptops on a four-year cycle. The company has performed a study and concluded that no other organization could use the laptops to gain a competitive edge, as they are general-purpose devices reaching the end of their useful life for the requirements of the company. The organization has also taken the steps of removing all internal hard drives and other media containing data from the laptops. Which of the following would be the best means of disposal for the laptops?
   - **A.** Place them in storage in the hopes that they will be used again within the company
   - **B.** Contract with an outside company to destroy the laptops
   - **C.** Donate the laptops to a school or other charity
   - **D.** Dispose of them by placing them in the trash

2.5 ANSWERS

1. **A** The company would be at the acquisition phase of the asset life cycle. Since it has already created requirements for the new software and has decided not to develop software internally, the company is faced with purchasing or licensing the software from an outside source. The acquisition phase allows the organization to acquire software without expending efforts and resources toward development. The scenario does not mention how the organization wishes to dispose of its older software, and it is already past the implementation and maintenance phases with the older software package.

2. **C** The most cost-effective, efficient, and environmentally friendly method of disposal would be to donate the laptops to a school or other type of charity. Since they have already been stripped of the media that contain sensitive data, they can provide no competitive edge to anyone else, and schools and charities most likely have lesser computing requirements than the company. It may not be cost-effective to contract with a local company to destroy them, and storing them for later use within the company is impractical since they have already reached the end of their useful life. Placing them in the trash may not be cost-effective or efficient, as regulations likely exist regarding recycling or sanitary disposal of sensitive electronic components in the laptops.



Objective 2.6 Determine data security controls and compliance requirements

In this last objective for Domain 2, we will discuss the security controls implemented to ensure asset security, with a specific focus on data/information security. We will also discuss how compliance influences the selection of security controls and how security controls are scoped and tailored. Additionally, we will cover three key data protection methods specified in the CISSP objectives: digital rights management, data loss prevention, and the cloud access security broker.

Data Security and Compliance

The controls for securing and protecting information cannot just be randomly picked. Many factors go into control selection, some of which include the state the data is in, the type of data, and the governance requirements for the data. Data states influence which controls are used, and how controls are implemented to meet the specific needs of the data that requires protecting.

Data States

Raw data and, when placed in context, information are categorized as always being in one of three states: at rest, in use (or in process), and in transit. Many of the security controls that are selected to protect data are focused on one or more of these three states, which we will discuss next.

Data in Transit

Data in transit is simply data that is being actively sent across a transmission media such as cabling or wireless. Data in this state is vulnerable to eavesdropping and interception if it is not properly protected. Controls traditionally used to protect data while in transit include strong authentication and transmission encryption. Examples of protocols and other mechanisms

used to protect data during transmission include Transport Layer Security (TLS), Secure Shell (SSH), hardware encryption devices, and so on.

Data at Rest

Data at rest is data that is in permanent storage, such as data that resides on a system's hard drive. It is data that is not being actively transmitted or received, nor is it actively being processed by the system. In this state, data is subject to unauthorized access, copying, deletion, or removal. Controls typically used to protect data while at rest include access controls such as permissions, strong encryption, and strong authentication.

Data in Use

Data in use is data that is in a temporary or ephemeral state, characterized by data being actively processed by a device's CPU and residing in volatile or temporary memory, such as RAM for example. During this state of being in use, data is transferred across the system bus to interact with various software and hardware components, including the CPU and memory, as well as the applications that use the data. Data in this state is destined to be read, transformed or modified, and returned back to storage.

This data state has been a traditional weakness in data and system security, in that data cannot always be fully protected by controls such as encryption. For example, different software and hardware components may not work with selected security controls because they need direct, unrestricted access to data. Only recently have efforts been made to protect data while it is being actively processed, and typically the controls that are used to protect it are now included in system hardware and software components. Advances in hardware and software have facilitated implementation of controls such as encrypting data between the storage device and the CPU on the system bus, memory protection, process isolation, and other security measures.



EXAM TIP Data exists in three recognized "states": at rest (in storage), in transit (being transmitted), and in use (being actively accessed and processed in a device's CPU and RAM).

Control Standards Selection

Selecting controls to protect information and other assets is not an arbitrary decision. Organizations use control standards such as the ISO/IEC 27000 series, the Center for Internet Security (CIS) Critical Security Controls (CIS Controls), or the National Institute of Standards and Technology (NIST) security controls (NIST SP 800-53, Revision 5, *Security and Privacy Controls for Information Systems and Organizations*) based upon several factors, including the data type (e.g., financial data, healthcare data, personal data, etc.), the level of protection needed (based on the sensitivity and criticality of the data), and, most importantly, compliance with governance. Very often, governance is the deciding factor, since some sensitive data types, such as credit card information, must be protected according to a specific control set.

Beyond governance, an organization may choose control standards because of the industry they are involved in or, data protection needs, or its management culture. For example, defense contractors providing services to the U.S Department of Defense typically follow the NIST controls by default because the U.S. federal government requires contractors to comply with those controls.

Scoping and Tailoring Data Security Controls

Regardless of the control standard that an organization chooses to protect its sensitive data, there are several factors that go into determining which controls the organization uses, and to what depth the organization implements the controls to protect its information. Controls are not simply implemented in a "one-size-fits-all" fashion; the controls selected, as well as their level of implementation, can be scoped and tailored appropriately to the organization's information protection requirements.

Considerations in scoping and tailoring information security controls include factors we discussed in Objective 1.10 when we talked about risk management, including the criticality and sensitivity of the data. The more critical or sensitive the data, the more stringent the controls that must be put into place to protect it. Another issue is the overall value of the data or asset to the organization. The organization conducts a cost/benefit analysis to determine if the value of the asset to the organization is far less than the cost to implement and maintain a control to protect it.

Governance also figures into this scoping and tailoring effort. Some laws and regulations require a specific level of control implementation, in addition to the basic standard they impose. For example, financial regulations require that personal financial information be encrypted, using a specific encryption strength or algorithm.

In any event, selecting controls and tailoring them the right way to protect data in critical assets at the level the organization requires is an exercise in risk management, as we discussed in Objective 1.10.

Data Protection Methods

This exam study guide covers a variety of data protection methods, but exam objective 2.6 identifies three specific methods that you should be familiar with for the exam, as described in this section. For this objective we will focus on digital rights management (DRM), data loss prevention (DLP), and cloud access security broker (CASB).

Digital Rights Management

Digital rights management (DRM) refers to various technologies that help to restrict access to data so that it cannot be copied or modified by an unauthorized user. DRM most often is implemented to protect copyrighted material, such as music and movies, but it is also used to protect proprietary or sensitive information that an organization produces, such as trade secrets. DRM is regularly implemented in software, and there are some hardware-enabled DRM solutions, such as DVD and Blu-ray players, that won't play copied or unencrypted discs.

Software implementations include protections like watermarking, file or software encryption, licensing limitations, and so on. A popular form of DRM is *steganography*, which involves concealment of embedded information in computer files. Steganography can be used in a malicious way to hide malware in otherwise innocuous-looking files, but in the DRM arena, it is used to insert digital watermarking into files to protect intellectual property.

Data Loss Prevention

Data loss prevention (DLP) also refers to a wide variety of technologies. In this case, they are used to prevent sensitive data from leaving the organization's infrastructure and control. Normally DLP is applied judiciously to only the most critical or sensitive data in an organization; using DLP to control every single piece of data that exits the network would be counterproductive and cost-prohibitive. DLP works primarily by assigning labels or metadata to files and other data elements. This metadata indicates the sensitivity of the data and serves as a marker so that devices and applications will not allow it to leave the network through e-mail, file copying, and so on.

The two main types of DLP are network DLP (NDLP) and endpoint DLP (EDLP).

- *Network DLP* is deployed as various network-enabled applications and devices that scan data leaving the network over transmission media and prevent data identified as sensitive from exiting over e-mail, HTTP connections, file transfers, and so on. NDLP can even intercept secure traffic leaving the company network, such as encrypted traffic, and decrypt that traffic to ensure that no sensitive data is leaving the company network.
- *Endpoint DLP* focuses on the end-user device, such as a workstation, and takes steps to ensure that data remains encrypted, is only accessible by authorized users, and cannot be copied from the device, such as through the use of USB and other removable media. EDLP prevents data from being copied from a centralized network storage location, which may even be the user's workstation. More often than not, hybrid solutions featuring both of these types of DLP are typically implemented within an organization.

A necessary requirement for DLP to work in an organization is that data must be "inventoried" by type, criticality, and sensitivity. Also, data flows, data processing, and storage locations must be discovered and documented throughout the infrastructure. DLP can control the flow of data within and outside of the network, provided the DLP solution understands where data resides and where it is supposed to flow.

Cloud Access Security Broker

A cloud access security broker (CASB) is a solution that is being used more often as organizations migrate data, applications, and even infrastructure to cloud providers. A CASB is a combination of software and hardware solutions that exist to mediate connections, data access, and data transfer to and from the organization's cloud storage and infrastructure. This mediation of access includes controls to implement strong authentication and encryption for both data

at rest and data in transit. The CASB provides visibility into what users are doing in relation to the organization's cloud services and applies security policies and controls for the data and applications in use.

The CASB uses two primary methods: proxying, with the use of a CASB appliance that sits in the data path between the organization and cloud assets, and application programming interfaces (APIs), which are integrated with cloud-enabled applications. While the organization may maintain some level of control over how the CASB solution controls access to the organization's data, the cloud service provider itself is usually responsible for maintaining the technology from functional and performance perspectives.

The primary advantage of using a CASB is centralized control over access to the organization's cloud-based resources, which provides easy, policy-based management. The disadvantages of using a CASB depend on the implementation method. Proxy-based CASBs must intercept traffic, including traffic that is encrypted by users. This means that encrypted traffic must be decrypted first and examined for security issues, such as data exfiltration, before it is re-encrypted and sent out to the cloud. This can introduce a variety of technical issues, as well as privacy issues, not to mention latency in the network connection. CASB APIs have their own vulnerabilities, which include ordinary software programming vulnerabilities that may be present in faulty code.

REVIEW

Objective 2.6: Determine data security controls and compliance requirements In this final objective of Domain 2 we closed out our discussion of asset security. We discussed various aspects of data security and compliance, including the three typical states of data. We covered the definitions and typical requirements for protection of data in transit, data at rest, and data in use. We also discussed the process for an organization to select security control standards used to protect assets. As part of this discussion, we covered key items an organization must consider when it is scoping and tailoring controls to protect information and other assets. We also covered three key data protection methods required by the exam objectives: DRM, DLP, and CASB.

2.6 QUESTIONS

1. You are a cybersecurity analyst in your company and have been tasked to explore which control standard the company will adopt to protect sensitive information. Your organization processes and stores a variety of sensitive data, including individual personal and financial information for your customers Which of the following is likely the most important factor in selecting a control standard?
   - **A.** Cost
   - **B.** Governance
   - **C.** Data sensitivity
   - **D.** Data criticality

2. Gary is a senior cybersecurity engineer for his company and has been tasked to implement a DLP solution. One of the solution requirements is that secure traffic leaving the company network, such as traffic encrypted with TLS, be intercepted and decrypted to ensure that no sensitive data is being exfiltrated. Which of the following must be included in the company's solution to meet that particular requirement?
   A. Network DLP
   B. DRM
   C. Endpoint DLP
   D. CASB

2.6 ANSWERS

1. **B** Since the organization processes sensitive data that may include personal or financial data of its customers, governance is likely the most important factor to consider when selecting a control standard, since personal and financial data are protected by various laws and regulations. These laws and regulations may require specific control standards. The other factors mentioned are important, but the organization has some leeway when considering those factors.
2. **A** Network data loss prevention (NDLP) must be included in the company's solution, since it is used specifically to protect against data being exfiltrated from an organization's network. Network DLP can be used to intercept encrypted traffic, decrypt and analyze the traffic, and then re-encrypt it before it exits the network. Digital rights management (DRM) and endpoint DLP (EDLP) should also be used as part of a layered solution, but they do not intercept and decrypt network traffic for analysis. A cloud access security broker (CASB) is only required for mediating access control for cloud-based assets.



Security Architecture and Engineering

Domain Objectives

- **3.1** Research, implement, and manage engineering processes using secure design principles.
- **3.2** Understand the fundamental concepts of security models (e.g., Biba, Star Model, Bell-LaPadula).
- **3.3** Select controls based upon systems security requirements.
- **3.4** Understand security capabilities of Information Systems (IS) (e.g., memory protection, Trusted Platform Module (TPM), encryption/decryption).
- **3.5** Assess and mitigate the vulnerabilities of security architectures, designs, and solution elements.
- **3.6** Select and determine cryptographic solutions.
- **3.7** Understand methods of cryptanalytic attacks.
- **3.8** Apply security principles to site and facility design.
- **3.9** Design site and facility security controls.

Domain 3, "Security Architecture and Engineering," is one of the largest and possibly most difficult domains to understand and remember for the CISSP exam. Domain 3 comprises approximately 13 percent of the exam questions. We're going to cover its nine objectives, which address secure design principles, security models, and selection of controls based on security requirements. We will also discuss security capabilities and mechanisms in information systems and go over how to assess and mitigate vulnerabilities that come with the various security architectures, designs, and solutions. Then we will cover the objectives that focus on cryptography, examining the basic concepts of the various cryptographic methods and understanding the attacks that target them. Finally, we will look at physical security, reviewing the security principles of site and facility design and the controls that are implemented within them.



Objective 3.1 Research, implement, and manage engineering processes using secure design principles

In this objective we will begin our study of security architecture and engineering by looking at secure design principles that are used throughout the process of creating secure systems. The scope of security architecture and engineering encompasses the entire systems and software life cycle, but it all begins with understanding fundamental security principles before the first design is created or the first component is connected. We have already touched upon a few of these principles in the previous domains, and in this objective (as well as throughout the remainder of the book) we will explore them in more depth.

Threat Modeling

Recall that we discussed the principles and processes of threat modeling back in Objective 1.11. Here, we will discuss them in the context of secure design. In addition to threat modeling as a process that should be performed on a continual basis throughout the infrastructure, threat modeling should also be considered during the architecture and engineering design phases of a system life cycle. As a reminder, threat modeling is the process of describing detailed threats, events, and their specific impacts on organizational assets. In the context of threat modeling as a secure design principle, cybersecurity architects and engineers should consider determining specific threat actors and events and how they will exploit a range of vulnerabilities that are inherent to a system or application that is being designed.

Least Privilege

As previously introduced in Objective 1.2, the principle of least privilege states that an entity (most commonly a user) should only have the minimum level of rights, privileges, and

permissions required to perform their job duties, and no more. The principle of least privilege is accomplished by strict review of an individual's access requirements and comparing them to what that individual requires to perform their job functions. The principle of least privilege should be practiced in all aspects of information security, to include system and data access (including physical), privileged use, and assignment of roles and responsibilities. This secure design principle should be used whenever a system or its individual components are bought or built and connected together to include mechanisms that restrict privileges by default.

Defense in Depth

Defense in depth means designing and implementing a multilayer approach to securing assets. The theory is that if multiple layers of protection are applied to an asset or organization, the asset or organization will still be protected in the event one of those layers fails. An example of this defense-in-depth strategy is that of a network that is protected by multiple levels at ingress/egress points. Firewalls and other security devices may protect the perimeter from most bad traffic entering the internal network, but other access controls, such as resource permissions, strong encryption, and authentication mechanisms are used to further limit access to the inside of the network from the outside world. Layers of defenses do not all have to be technical in nature; administrative controls in the form of policies and physical controls in the form of secured processing areas are also used to add depth to these layers.

Secure Defaults

Security and functionality are often at odds with each other. Generally, the more functional a system or application is for its users, the less secure it is, and vice versa. As a result, many applications and systems are configured to favor functionality by permitting a wider range of actions that users and processes can perform. The focus on functionality led organizations to configure systems to "fail open" or have security controls disabled by default. The principle of *secure defaults* means that out-of-the-box security controls should be secure, rather than open. For instance, when older operating systems were initially installed, cybersecurity professionals had to "lock down" the systems to make them more secure, since by default the systems were intended to be more functional than secure. This meant changing default blank or simple passwords to more complex ones, implementing encryption and strong authentication mechanisms, and so on. A system that follows the principle of secure defaults already has these controls configured in a secure manner upon installation, hence its default state.

Fail Securely

Another application of the secure default principle is when a control in a system or application fails due to an error, disruption of service, loss of power, or other issue, the system fails in a secure manner. For example, if a system detects that it is being attacked and its resources, such

as memory or CPU power, are degraded, the system will automatically secure itself, preventing access to data and critical system components. The related term for the secure default principle during a failure is "fail secure," which contrasts with the term "fail safe."

Although in some controls the desired behavior is to fail to a secure state, other controls must fail to a safe mode of operation in order to protect human safety, prevent equipment damage and data loss, and so on. An example of the fail-safe principle would be when a fire breaks out in a data center, the exit doors fail to an open state, rather than a secured or locked state, to allow personnel to safely evacuate the area.



EXAM TIP Secure default is associated with the term "fail closed," which means that in the event of an emergency or crisis, security mechanisms are turned on. Contrast this to the term "fail open," which can also be called "fail safe," and means that in the event of a crisis, security controls are turned off. There are situations where either of those conditions could be valid responses to an emergency situation or crisis.

Separation of Duties

The principle of separation of duties (SoD) is a very basic concept in information security: one individual or entity should not be able to perform multiple sensitive tasks that could result in the compromise of systems or information. The SoD principle states that tasks or duties should be separated among different entities, which provides a system of checks and balances. For example, in a bank, tellers have duties that may allow them to access large sums of money. However, to actually transfer or withdraw money from an account requires the signature of another individual, typically a supervisor, in order for that transaction to occur. So, one individual cannot simply transfer money to their own account or make a large withdrawal for someone else without a separate approval.

The same applies in information security. A classic example is that of an individual with administrative privileges, who may perform sensitive tasks but is not allowed to also audit those tasks. Another individual assigned to access and review audit records would be able to check the actions of the administrator. The administrator would not be allowed to access the audit records since there's a possibility that the administrator could alter or delete those records to cover up their misdeeds or errors. The use of defined roles to assign critical tasks to different groups of users is one way to practically implement the principle of separation of duties.

A related principle is that of multiperson control, also sometimes referred to as "M of N" or two-person control. This principle states two or more people are required to perform a complete task, such as accessing highly sensitive pieces of information (for example, the enterprise administrator password). This principle helps to eliminate the possibility of a single person having access to sensitive information or systems and performing a malicious act.

Keep It Simple

Keep it simple is a design principle that means that architectures and engineering designs should be kept as simple as possible. The more complex a system or mechanism, the more likely it is to have inherent weaknesses that will not be discovered or security mechanisms that may be circumvented. Additionally, the more complex a system, the more difficult it is to understand, configure, and document. Security architects and engineers must avoid the temptation to unnecessarily overcomplicate a security mechanism or process.

Zero Trust

The principle of zero trust states that no entity in the infrastructure automatically trusts any other entity; that is to say, each entity must always reestablish trust with another one. Under this principle, an entity is considered hostile until proven otherwise. For example, hosts in a network should always have to mutually authenticate with each other to verify the other host's identity and access level, even if they have performed this action before. Additionally, the principle ensures that even when trust is established, it is kept as minimal as possible. Trust between entities is very defined and discrete—not every single action, process, or component in a trusted entity is also considered trusted. This principle mitigates the possibility that, if the host or other entity is compromised since the last time they established trust, the systems are able to communicate or exchange data.

Mutual authentication, periodic reauthentication, and replay prevention are three key security measures that can help establish and support the zero-trust principle. Since widespread use of the zero-trust principle throughout an infrastructure could hamper data communication and exchange, implementation is usually confined to only extremely sensitive assets, such as sensitive databases or servers.

Privacy by Design

Privacy by design means that considerations for individual privacy are built into the system and applications from the beginning, including as part of the initial requirements for the system and continuing into design, architecture, and implementation. Remember that privacy and security are not the same thing. Security seeks to protect information from unauthorized disclosure, modification, or denial to authorized entities. Privacy seeks to ensure that the control of certain types of information is kept by the individual subject of that information. Privacy controls built into systems and applications ensure that privacy data types are marked with the appropriate metadata and protected from unauthorized access or transfer to an unauthorized entity.

Trust But Verify

Trust but verify is the next logical extension of the zero-trust model. In the zero-trust model, no entity is automatically trusted. The trust but verify principle goes one step further by requiring

that once trust is established, it is periodically reestablished and verified, in the event one entity becomes compromised. Auditing is also an important part of the verification process between two trusted entities; critical or sensitive transactions are audited and monitored closely to ensure that the trust is warranted, is within accepted baselines, and has not been broken.

Shared Responsibility

Shared responsibility is a model that applies when more than one entity is responsible and accountable for protecting systems and information. Each entity has its own prescribed tasks and activities focused on protecting systems and data. These responsibilities are formally established in an agreement or contract and appropriately documented. The best example of a shared responsibility model is that of a cloud service provider and its client, who are each responsible for certain aspects of securing systems and information as well as their access. The organization may maintain responsibilities on its end for initially provisioning user access to systems and data, while the cloud provider is responsible for physical and technical security protections for those systems and data.

REVIEW

Objective 3.1: Research, implement, and manage engineering processes using secure design principles For the first objective of Domain 3, we discussed foundational principles that are critical during security architecture and engineering design activities.

- *Threat modeling* should occur not only as a routine process but also before designing and implementing systems and applications, since security controls can be designed to counter those threats before the system is implemented.
- The *principle of least privilege* states that entities should only have the required rights, permissions, and privileges to perform their job function, and no more.
- *Defense in depth* is a multilayer approach to designing and implementing security controls to protect assets in the organization. It is employed such that in the event that one or more security controls fail, others can continue to protect sensitive assets.
- *Secure default* is the principle that states that when a system is first implemented or installed, its default configuration is a secure state.
- *Fail secure* means that a system should also fail to a secure state when it is unexpectedly halted, interrupted, or degraded.
- *Separation of duties* means that critical or sensitive tasks should not all fall onto one individual; these tasks should be separated amongst different individuals to prevent one person from being able to cause serious damage to systems or information.

- *Keep it simple* means that the more complex the system or application, the less secure it is and more difficult to understand and document.
- *Zero trust* means that two or more entities in an infrastructure do not start out by trusting each other. Trust must first be established and then periodically reestablished.
- *Trust but verify* is a principle that means that once trust is established, it must be periodically reestablished and verified as still current and necessary, in the event one entity becomes compromised.
- *Privacy by design* is the principle that privacy considerations should be included in the initial requirements, design, and architecture for systems, applications, and processes, so that individual privacy can be protected.
- *Shared responsibility* is the principle that means that two or more entities share responsibility and accountability for securing systems and data. This shared responsibility should be formally agreed upon and documented.

3.1 QUESTIONS

1. You are a cybersecurity administrator in a large organization. IT administrators frequently perform many sensitive tasks, to include adding user accounts and granting sensitive data access to those users. To ensure that these administrators do not engage in illegal acts or policy violations, you must frequently check audit logs to make sure they are fulfilling their responsibilities and are accountable for their actions. Which of the following principles is employed here to ensure that IT administrators cannot audit their own actions?
   A. Principle of least privilege
   B. Trust but verify
   C. Separation of duties
   D. Zero trust
2. Your company has just implemented a cutting-edge data loss prevention (DLP) system, which is installed on all workstations, servers, and network devices. However, the documentation for the solution is not clear regarding how data is marked and transits throughout the infrastructure. There have been several reported instances of data still making it outside of the network during tests of the solution, due to multiple possible storage areas, transmission paths, and conflicting data categorization. Which of the following principles is likely being violated in the secure design of the solution, which is allowing sensitive data to leave the infrastructure?
   A. Keep it simple
   B. Shared responsibility
   C. Secure defaults
   D. Fail secure

3.1 ANSWERS

1. **C** The principle of separation of duties is employed here to ensure that critical tasks, such as auditing administrative actions, are not performed by the people whose activities are being audited, which in this case are the IT administrators. Auditors are responsible for auditing the actions of the IT administrators, and these two tasks are kept separate to ensure that unauthorized actions do not occur.
2. **A** In this scenario, the principle of keep it simple is likely being violated, since the solution may be configured in an overly complex manner and is allowing data to traverse multiple uncontrolled paths. A further indication that this principle is not being properly employed is the lack of clear documentation for the solution, as indicated in the scenario.



Objective 3.2 Understand the fundamental concepts of security models (e.g., Biba, Star Model, Bell-LaPadula)

A *security model* (sometimes also called an *access control model*) is a mathematical representation of how systems and data are accessed by entities, such as users, processes, and other systems. In this objective we will examine security models and explain how they allow access to data based on a variety of criteria, including security clearance, need-to-know, and job roles.

Security Models

Security models propose how to allow entities controlled access to information and systems, maintaining confidentiality and/or integrity, two key goals of security. *Security levels* are different sensitivity levels assigned to information systems. These levels come directly from data sensitivity policies or classification schemes. Models that approach access control from the confidentiality perspective don't allow an entity at one security level to read or otherwise access information that resides at a different security level. Models that approach access control from the integrity perspective do not allow subjects to write to or change information at a given security level. Most of the security models we will discuss are focused on those two goals, confidentiality and integrity.

Security models are most often associated with the *mandatory access control (MAC)* paradigm of access control, which is used by administrators to enforce highly restrictive access controls on users and other entities (called *subjects*) and their interactions with systems and information, often referred to as *objects*.

Cross-Reference

Mandatory access control (MAC), discretionary access control (DAC), role-based access control (RBAC), and other access control models are discussed in Objective 5.4.

Terms and Concepts

Access to information in mandatory access control models is based upon the three important concepts of security clearance, management approval, and *need-to-know*. A security clearance verifies that an individual person is "cleared" for information at a given security level, such as the U.S. military classification levels of "Secret" and "Top Secret." A person with a Secret security clearance is not allowed to access information at the Top Secret level due to the difference in classification level permissions. Access to a given system or set of information is denied when it has a security level higher than that of the subject desiring access. Need-to-know also restricts access even if the subject has equal or higher security clearance than the system or data.

As mentioned earlier, the type of access, for our purposes here, is the ability to read information from or write information to that security level. If a subject is not authorized to "read up" to a higher (more restrictive) security level, they cannot access information or systems at that level. If a subject is not authorized to write to a different security level (regardless of whether or not they can read it), then they cannot modify information at that level. Based on these restrictions, it would seem that if a subject has a clearance to access information at a given level, they would automatically be allowed to access information at any lower (less restrictive) level. This is not necessarily the case, however, because the additional requirements of need-to-know and management approval must be applied.

Even if a subject has the clearance to access information at a higher (or even the same) security level, if they don't have a valid need-to-know related to their job position, then they should not be allowed access to that information, regardless of whether it resides at a lower security level or not. Additionally, management must formally approve the access; simply having the clearance and need-to-know does not automatically grant information access. Figure 3.2-1 shows how the concepts of confidentiality and integrity are implemented in these examples.



NOTE We often look at multilevel security as being between levels that are higher (more restrictive) or lower (less restrictive) than each other, but this is not necessarily the case. Information can also be compartmented (even at the same "level") but still separated and restricted in terms of access controls. This is where "need-to-know" comes into play the most.

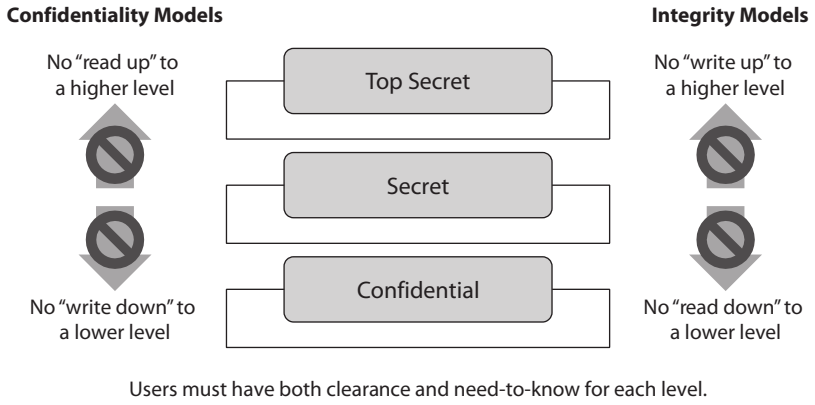


FIGURE 3.2-1 Ensuring confidentiality and integrity in mandatory access control

System States and Processing Modes

Before we discuss the various confidentiality and integrity models, it's helpful to understand the different "states" and "modes" that a system operates in for information categorized at different security levels. In a *single-state* system, the processor and system resources may handle only one security level at a time. A single-state system should not handle multiple security levels concurrently, but could switch between them using security policy mechanisms. A *multi-state* system, on the other hand, can handle several security levels concurrently, because it has specialized security mechanisms built in. Different processing modes are defined based on the security levels processed and the access control over the user. There are four generally recognized processing modes you should be familiar with for the exam, presented next in order of most restrictive to least restrictive.

Dedicated Security Mode

Dedicated security mode is a single-state type of system, meaning that only one security level is processed on it. However, the access requirements to interact with information are restrictive. Essentially, users that access the security level must have

- A security clearance allowing them to access *all* information processed on the system
- Approval from management to access *all* information processed by the system
- A valid need-to-know, related to their job position, for *all* information processed on the system

System-High Security Mode

For system-high security mode, the user does not have to possess a valid need-to-know for *all* of the information residing on the system but must have the need-to-know for at least *some* of it. Additionally, the user must have

- A security clearance allowing them to access *all* information processed on the system
- Approval from management and signed nondisclosure agreements (NDAs) for *all* information processed by the system

Compartmented Security Mode

Compartmented security mode is slightly less restrictive in terms of allowing user access to the system. For compartmented security mode, users must meet the following requirements:

- Level of security clearance allowing them to access *all* information processed on the system
- Approval from management to access *at least some* information on the system and signed NDAs for all information on the system
- Valid need-to-know for *at least* the information they will access

Multilevel Security Mode

Finally, multilevel security mode may have multiple levels of sensitive information processed on the system. Users are not always required to have the security clearance necessary for all of the information on the system. However, users must have

- Security clearance at *least* equal to the level of information they will access
- Approval from management for *any* information they will access and signed NDAs for all information on the system
- A valid need-to-know for *at least* the information they will have access to

Figure 3.2-2 illustrates the concepts associated with multilevel security.



EXAM TIP Understanding the security states and processing modes is key to understanding the confidentiality and integrity models we discuss for this objective, although you may not be asked specifically about the states or processing modes on the exam.

Mode	Clearance	Access Approval	Need-to-know
Dedicated	All information processed on system	All information processed on system	All information processed on system
System-high	All information processed on system	All information processed on system	Some information processed on system
Compartmented	All information processed on system	Any information processed on system	All information they will have access to on the system
Multilevel	All information they will have access to on the system	All information they will have access to on the system	All information they will have access to on the system

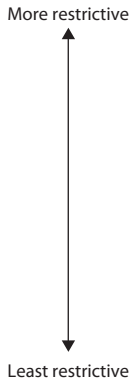


FIGURE 3.2-2 Multilevel security states and modes

Confidentiality Models

As mentioned, security models target either the confidentiality aspect or the integrity aspect of security. Confidentiality models, discussed here, seek to strictly control access to information, namely the ability to read information at specific security levels. Confidentiality models, however, do not consider other aspects, such as integrity, so they do not address the potential for unauthorized modification of information by a subject that may be able to "write up" to the security level, even if they cannot read information at that level.

Bell-LaPadula

The most common example of a confidentiality access control model is the Bell-LaPadula model. It only addresses confidentiality, not integrity, and uses three main rules to enforce access control:

- **Simple security rule** A subject at a given security level is not allowed to read data that resides at a higher or more restrictive security level. This is commonly called the "no read up" rule, since the subject cannot read information at a higher classification or sensitivity level.
- ***-property rule (called the star property rule)** A subject at a given security level cannot write information to a lower security level; in other words, the subject cannot "write down," which would transfer data of a higher sensitivity to a level with lower sensitivity requirements.
- **Strong star property rule** Any subject that has both read and write capabilities for a given security level can only perform both of those functions at the *same* security level—nothing higher and nothing lower. For a subject to be able to both read and write to an object, the subject's clearance level must be equal to that of the object's classification or sensitivity.

Figure 3.2-1, shown earlier, demonstrates how the simple security and star property rules function in a confidentiality model.



EXAM TIP The CISSP exam objectives specifically mention the "Star Model," but this is a reference to the star property and star integrity rules (often referred to in shorthand as "* property" and "* integrity") found in confidentiality and integrity models.

Integrity Models

As mentioned, some mandatory access control models only address integrity, with the goals of ensuring that data is not modified by subjects who are not allowed to do so and ensuring that data is not written to different classification or security levels. Two popular integrity models are Biba and Clark-Wilson, although there are many others as well.

Biba

The Biba model uses integrity levels (instead of security levels) to prevent data at any integrity level from flowing to a different integrity level. Like Bell-LaPadula, Biba uses three primary rules that affect reading and writing to different security levels and uses them to enforce integrity protection instead of confidentiality. These levels are also illustrated in Figure 3.2-1.

- **Simple integrity rule** A subject cannot read data from a lower integrity level (this is also called "no read down").
- ***-integrity rule** A subject cannot write data to an object residing at a higher integrity level (called "no write up").
- **Invocation rule** A subject at one integrity level cannot request or invoke a service from a higher integrity level.



EXAM TIP Although Bell-LaPadula (confidentiality model) and Biba (integrity model) address different aspects of security, both are called *information flow models* since they control how information flows between security levels. For the exam, remember the rules that apply to each. If the word "simple" is used, the rule addresses *reading* information. If the rule uses the * symbol or "star" in its name, the rule applies to *writing* information.

Clark-Wilson

The Clark-Wilson model is also an integrity model but was developed after Biba and uses a different approach to protect information integrity. It uses a technique called *well-formed transactions*, along with strictly defined separation of duties. A well-formed transaction is a series of

operations that transforms data from one state to another, while keeping the data consistent. The consistency factor ensures that data is not degraded and preserves its integrity. Separation of duties between processes and subjects ensures that only valid subjects can transform or change data.

Other Access Control Models

We covered two of the key security models, confidentiality and integrity, involved in mandatory access control, and there are several other models that perform some of the same functions, using different techniques. In reality, neither of these security models is very practical in the real world; each has its place in very specific contexts. For example, a classified military system or even a secure healthcare system may use different kinds of access control models to protect information at different stages of processing, so they do not rely only on one type of access control model. More often than not, you will see these security models used together under different circumstances.

Although we will not discuss them in depth here, additional models that you may need to be familiar with for the purposes of the exam include

- **Noninterference model** This multilevel security model requires that commands or activities performed at one security level should not be seen by or affect subjects and objects at a different security level.
- **Brewer and Nash model** Also called the Chinese Wall model, this model is used to protect against conflicts of interest by stating that a subject can write to an object if and only if the subject cannot read another object from a different set of data.

REVIEW

Objective 3.2: Understand the fundamental concepts of security models (e.g., Biba, Star Model, Bell-LaPadula) In this objective we examined the fundamentals of security models, which typically follow the paradigm of mandatory access control. We examined models that use two different approaches to access control: confidentiality models and integrity models. Confidentiality models are concerned with strictly controlling the ability of a subject to "read" or access information at a given security level. The main confidentiality example is the Bell-LaPadula model, which uses rules to inhibit the ability of a subject to "read up" and "write down," to prevent unauthorized data access. Integrity models rigorously restrict the ability of subjects to write to or modify information at given security levels. We also discussed system states and processing modes, which include dedicated security mode, system-high security mode, compartmented security mode, and multilevel security mode.

3.2 QUESTIONS

1. If a user has a security clearance equal to *all* information processed on the system, but is only approved by management for *some* information and only has a need-to-know for that specific information, which of the following is the security mode for the system?
   - **A.** Compartmented security mode
   - **B.** Dedicated security mode
   - **C.** Multilevel security mode
   - **D.** System-high security mode
2. Mandatory access control security models address which of the following two goals of security?
   - **A.** Confidentiality and availability
   - **B.** Availability and integrity
   - **C.** Confidentiality and integrity
   - **D.** Confidentiality and nonrepudiation
3. Your supervisor has granted you rights to access a system that processes highly sensitive information. In order to actually access the system, you must have a security clearance equal to *all* information processed on the system, management approval for *all* information processed on the system, and need-to-know for *all* information processed on the system. In which of the following security modes does the system operate?
   - **A.** Multilevel security mode
   - **B.** System-high security mode
   - **C.** Dedicated security mode
   - **D.** Compartmented security mode
4. You are a cybersecurity engineer helping to design a system that uses mandatory access control. The goal of the system is to preserve information integrity. Which of the following rules would be used to prevent a subject from writing data to an object that resides at a higher integrity level?
   - **A.** Strong security rule
   - **B.** *-security rule
   - **C.** Simple integrity rule
   - **D.** *-integrity rule

3.2 ANSWERS

1. **A** Compartmented security mode means that the user must have a security clearance equal to *all* information processed on the system, regardless of management approval to access the information or need-to-know. Additionally, the user must have specific management approval and need-to-know for at least *some* of the information processed on the system.
2. **C** Mandatory access control models address the confidentiality and integrity goals of security.
3. **C** Since the user must meet all three requirements for *all* information processed on the system (security clearance, management approval, and need-to-know), the system operates in dedicated security mode. This indicates a single-state system since it is using only one security level.
4. **D** The *-integrity rule states that a subject cannot write data to an object at a higher integrity level (called "no write up").



Objective 3.3 Select controls based upon systems security requirements

In this objective we will discuss how to select security controls for implementation based upon systems security requirements. We discussed some of these requirements in Domain 1. In this objective we will review them in the context of the entire system's life cycle and, more specifically, how to select security controls during the early stages of that life cycle.

Selecting Security Controls

As you may remember from Objective 1.10, we select security controls to implement based on a variety of criteria. In addition to cost, we consider the protection needs of the asset, the risk the control reduces, and how well the control complies with governance requirements. We should reconsider implementing a control if it either costs more than the asset is worth or does not provide sufficient protection, risk reduction, or compliance to balance its costs.

Controls should not be selected for implementation after the system or application is already installed and working; the control selection process actually takes place before the system is built or bought. To identify the type of security controls to implement prior to installation and operation of the system, we must review the full requirements of the system. This review takes place during the requirements phase of the software/system development life cycle (SDLC) and includes functional requirements (what it must do), performance requirements (how well it must do it), and, of course, security requirements. Also, remember that technical controls

are not the only ones to consider; administrative and physical controls go hand-in-hand with technical controls and should be considered as well. While there are many variables to consider when selecting which security controls to implement, we'll focus on several key requirements that you should understand for purposes of the exam.

Cross-Reference

We'll discuss the SDLC in greater detail in Objective 8.1.

Performance and Functional Requirements

Security requirements state what the system, software, or even a mechanism has to do; in other words, what function it fulfills. To meet security requirements, controls have to be selected based upon their functional requirements (what they do) and their ability to perform that function (how well they do it). Functional and performance requirements for controls should be established before a system or application is acquired (built or purchased). A control may be selected because it provides a desired security functionality that supports confidentiality, integrity, and/or availability. From a practical perspective, this functionality might be encryption or authentication.

We also have to consider how well the control performs its function. Requirements should always describe acceptable performance criteria for the system the control will support. Simply selecting a control that provides authentication services is not enough; it must meet specific standards, such as encryption strength, timing, interoperability with other systems, and so on. For example, choosing an older authentication mechanism that only supports basic username and password authentication or unencrypted protocols would likely not meet expected performance standards today. Functionality and performance requirements can be traced back to data protection and governance directives as well.

Data Protection Requirements

Data protection requirements can come from different sources, including governance. However, the two key factors in data protection are data criticality and sensitivity, as discussed in depth in Objective 2.1. Remember that *criticality* describes the importance of the information or systems to the organization in terms of its business processes and mission. *Sensitivity* refers to the information's level of confidentiality and relates to its vulnerability to unauthorized access or disclosure. These two requirements should be considered when selecting controls.

Controls must take into account the level of criticality of the system or the information that system processes and be implemented to protect that asset at that level. A business impact analysis (BIA), discussed in Objective 1.8, helps an organization determine the criticality of its business processes and identify the assets that support those processes. Using the BIA findings, an organization can choose and implement appropriate controls to protect critical processes.

Controls that protect critical data provide resiliency and availability of assets. Examples of these types of controls include backups, system redundancy and failover, and business continuity and disaster recovery plans.

As mentioned, sensitivity must also be considered when selecting and implementing controls. The more sensitive the systems and information, the stronger the control should be to protect their confidentiality. Controls designed to protect confidentiality include strong encryption, authentication, and other access control mechanisms, such as rights, permissions, and defined roles, and so on.

Governance Requirements

Beyond the organization's own determination of criticality and sensitivity, governance heavily influences control selection and implementation. Governance may levy specific mandatory requirements on the controls that must be in place to protect systems and information. For example, requirements may include a specific level of encryption strength or algorithm, or physical controls that call for inclusion of guards and video cameras in secure processing areas where sensitive data is accessed. Examples of governance that have defined control requirements include the Health Insurance Portability and Accountability Act (HIPAA), the Payment Card Industry Data Security Standard (PCI DSS), Sarbanes-Oxley (SOX), and various privacy regulations.

Interface Requirements

Interface requirements influence controls at a more detailed level and are quite important. These requirements dictate how systems and applications interact with other systems and applications, processes, and so on. Interface requirements should consider data exchange, formatting, and access control. Since sensitive data may traverse one or several different networks, it's important to look at the controls that sit between those networks or systems to ensure that they allow only the authorized level of data to traverse them, under very specific circumstances. For example, if a sensitive system is connected to systems of a lower sensitivity level, data must not be allowed to travel between systems unless it has been appropriately sanitized or processed so that it is the same sensitivity level of the destination system. A more specific example is a system in a hospital that contains protected healthcare information; if the system is connected to financial systems, certain information should be restricted from flowing to those systems, but other information, such as the patient's name and billing information, must be transferred to those systems. Controls at this level could be responsible for selectively redacting private health information that is not required for financial transactions.

Controls that are selected based upon interface requirements include secure protocols and services that move information between systems, strong encryption and authentication mechanisms, and network security devices, such as firewalls.

Risk Response Requirements

Risk response requirements that the controls must fulfill may be hard to nail down during the requirements phase. This is why a risk assessment and analysis must take place—before the system even exists. Risk assessments conducted during the requirements phase gather data about the threat environment and take into account existing security controls that may already be in place to protect an asset. Controls selected for implementation may be above and beyond those already in place but are necessary to bridge the gap between those controls and the potentially increased threats facing a new system.

A risk assessment and analysis should take into account all the other items we just discussed: governance, criticality and sensitivity factors, interconnection and interface requirements, as well as many other factors. Other factors that should be considered include the threat landscape (threat actors and known threats to the organization and its assets); potential vulnerabilities in the organization, its infrastructure, and in the asset (even before it is acquired); and the physical environment (facilities, location, and other environmental factors).



EXAM TIP Controls are considered and selected based upon several different factors, including functional and performance requirements, governance, the interfaces they will protect, and responses to risk analysis. However, the most important factor in selecting security controls is how well they protect systems and data.

Threat modeling contributes to the requirements phase risk assessment by developing a lot of the risk information for you; you will discover vulnerabilities and other particulars of risk as an incidental benefit to the threat modeling effort. The key to selecting controls based on risk response is weighing the existing controls against the threats that are identified for the asset and organization, and then closing the gap between the current security posture and the desired security state once the asset is in place.

Cross-Reference

Objective 1.11 provided in-depth coverage of threat modeling.

REVIEW

Objective 3.3: Select controls based upon systems security requirements In this objective we expanded our discussion of security control selection and discussed how these controls should be considered and selected even before the asset is acquired. This takes place during the requirements phase of the SDLC and includes considerations for functionality and performance of the control, data protection requirements, governance, interface requirements, and even risk response.

3.3 QUESTIONS

1. Your company is considering purchasing a new line-of-business application and integrating it with a legacy infrastructure. In considering additional controls for the new application, management has already taken into account governance, as well as how the controls must perform. However, it has not yet considered security controls affecting the interoperability of the application with other components. Which of the following requirements should management take into account when selecting security controls for the new application?
   - **A.** Functionality requirements
   - **B.** Risk response requirements
   - **C.** Interface requirements
   - **D.** Authentication requirements
2. Which of the following should be conducted during the requirements phase of the SDLC to adequately account for new threats, potential asset vulnerabilities, and other organizational factors that may affect the selection of additional security controls?
   - **A.** Business impact analysis
   - **B.** Risk assessment
   - **C.** Interface assessment
   - **D.** Controls assessment

3.3 ANSWERS

1. **C** Interface requirements address the interconnection of the application to other system components, as well as its interoperability with those legacy components. Security controls must consider how data is exchanged between those components in a secure manner, especially if the legacy components cannot use the same security mechanisms.
2. **B** During the requirements phase, a risk assessment should be conducted to ascertain the existing state of controls as well as any new or emerging threats that could present a problem for new assets. Additionally, the organizational security posture should be considered. A risk assessment will help determine new controls that would be available for risk response.



Objective 3.4 Understand security capabilities of Information Systems (IS) (e.g., memory protection, Trusted Platform Module (TPM), encryption/decryption)

In this objective we will explore some of the integrated security capabilities of information systems. These security solutions are built into the hardware and firmware, such as the Trusted Platform Module, the hardware security module, and the self-encrypting drive. We also will briefly discuss bus encryption, which is used to protect data while it is accessed and processed within the computing device. Additionally, we will examine security concepts and processes such as the trusted execution environment, processor security extensions, and atomic execution. Understanding each of these concepts is important to fully grasp how system security works at the lower hardware levels.

Information System Security Capabilities

While many types of information system security capabilities are available, we are going to limit the discussion here to the security capabilities of hardware and firmware that are specified or implied in exam objective 3.4. We will discuss other information system security capabilities throughout the book, including cryptographic capabilities in Objectives 3.6 and 3.7.

Hardware and Firmware System Security

Hardware and firmware security capabilities embedded in a system are designed to address two main issues: trust in the system's components; and protection of data at rest (residing in storage), in transit (actively being sent across a transmission media), and in use (actively being processed by a device's CPU and residing in volatile or temporary memory). Again, although there are several components that contribute to these processes, we're going to talk about the specific ones that you are required to know to meet the exam objective.

Trusted Platform Module

A Trusted Platform Module (TPM) is implemented as a hardware component installed on the main board of the computing device. Most often it is implemented as a computer chip. Its purpose is to carry out several security functions, including storing cryptographic keys and

digital certificates. It also performs operations such as encryption and hashing. TPMs are used in two important scenarios:

- **Binding a hard disk drive** This means that the hard drive is "keyed" through the use of encryption to work only on a particular system, which prevents the hard drive from being stolen and used in another system in order to gain access to its data.
- **Sealing** This is the process of encrypting the data for a system's specific hardware and software configuration and storing it on the TPM. This method is used to prevent tampering with hardware and software components to circumvent security mechanisms. If the drive or system is tampered with, the drive cannot be accessed.

TPMs use two different types of memory to store cryptographic keys: persistent memory and versatile memory. The type of memory used for each key and other security information depends on the purpose of the key. *Persistent memory* maintains its contents even when power is removed from the system. *Versatile memory* is dynamic and will lose its contents when power is turned off or lost, just as normal system RAM (volatile memory) does. The types of keys and other information stored in these memory areas include the following:

- **Endorsement key (EK)** This is the public/private key pair installed in the TPM when it is manufactured. This key pair cannot be modified and is used to verify the authenticity of the TPM. It is stored in persistent memory.
- **Storage root key (SRK)** This is the "master" key used to secure keys stored in the TPM. It is also stored in persistent memory.
- **Platform configuration registers (PCRs)** These are used to store cryptographic hashes of data and used to "seal" the system via the TPM. These are part of the versatile memory of the TPM.
- **Attestation identity keys (AIKs)** These keys are used to attest to the validity and integrity of the TPM chip itself to various service providers. Since these keys are linked to the TPM's identity when it is manufactured, they are also linked to the endorsement key. These keys are stored in the TPM's versatile memory.
- **Storage keys** These keys are used to encrypt the storage media of the system and are also located in versatile memory.

Hardware Security Module

A hardware security module (HSM) is almost identical in function to a TPM, with the difference being that a TPM is implemented as a chip on the main board of the computing device, whereas an HSM is a peripheral device and can be connected externally to systems that do not possess a TPM via an add-on card or even a USB connection.



EXAM TIP A TPM and HSM are almost identical and serve the same functions; the difference is that a TPM is built into the system's mainboard and an HSM is a peripheral device in the form of an add-on card or USB device.

Self-Encrypting Drive

A self-encrypting drive (SED), as its name suggests, is a self-contained hard disk that has encryption mechanisms built into the drive electronics; it does not require the TPM or HSM of a computing device. The key is stored within the drive itself and can be managed by a password chosen by the user. An SED can be moved between devices, provided they are compatible with the drive.

Bus Encryption

Bus encryption was developed to solve a potential issue that results when data must be decrypted from permanent storage before it is used by applications and hardware on a system. During that transition, the data is in use (active) and, if unencrypted, is vulnerable to being read by a malicious application and sent to a malicious entity. Bus encryption encrypts data before it is put on the system bus and ensures that data is encrypted even within the system while it is in use, except when being directly accessed by the CPU. However, bus encryption requires the use of a *cryptoprocessor*, which is a specialized chip built into the system to manage this process.

Secure Processing

There are several key characteristics of secure processing that you need to be familiar with for the CISSP exam, which include those managed by the hardware and firmware discussed in the previous section. Note that these are not the only secure processing mechanisms, but for objective 3.4, we will focus on the trusted execution environment, processor security extensions, and atomic execution.

Trusted Execution Environment

A trusted execution environment (TEE) is a secure, separated environment in the processor where applications can run and files can be open and protected from outside processes. Trusted execution environments are used widely in mobile devices (Apple refers to its implementation as a *secure enclave*). TEEs work by creating a trusted perimeter around the application processing space and controlling the way untrusted objects interact with the application and its data. The TEE has its own access to hardware resources, such as memory and CPU, that are unavailable to untrusted environments and applications.

You should note that a TEE protects several aspects of the processing environment; since the application runs in memory space that should not be corrupted by any other processes, this is referred to as *memory protection*. Memory protection is also achieved using a combination of other mechanisms as well, such as operating system and software application controls.



EXAM TIP A TEE is a protected processing environment; it is not a function of the TPM itself, but it does rely on the TPM to provide a chain of trust to the device hardware so the trusted software environment can be created.

Processor Security Extensions

Modern hardware built with security in mind can provide the needed security boundary to create a TEE. Modern CPUs come with *processor security extensions*, which are instructions specifically created to provide trusted environment support. Processor security extensions enable programmers to create special reserved regions in memory for sensitive processes, allowing these memory areas to contain encrypted data that is separated from all other processes. These memory areas can be dynamically decrypted by the CPU using the security extensions. This protects sensitive data from any malicious process that attempts to read it in plaintext. Processor security extensions also provide the ability to control interactions between applications and data in a TEE and untrusted applications and hardware that reside outside the secure processing perimeter.

Atomic Execution

Atomic execution is not so much a technology as it is a method of controlling how parts of applications run. It is an approach that prevents other, nonsecure processes from interfering with resources and data used by a protected process. To implement atomic execution, programmers leverage operating system libraries that invoke hardware protections during execution of specific code segments. A disadvantage of this, however, is that it can result in performance degradation for the system. An atomic operation is either fully executed or not performed at all.

Atomic execution is specifically designed to protect against timing attacks, which exploit the dependencies on sequence and timing in applications to execute multiple tasks and processes concurrently. These attacks are routinely called *time-of-check to time-of-use (TOC/TOU)* attacks. They attempt to interrupt the timing or sequencing of tasks that segments of code must complete. If an attacker can interrupt the sequence or timing after specific tasks, it can cause the application to fail, at best, or allow an attacker to read sensitive information, at worst. This type of attack is also referred to as an *asynchronous attack*.

REVIEW

Objective 3.4: Understand security capabilities of Information Systems (IS) (e.g., memory protection, Trusted Platform Module (TPM), encryption/decryption) In this objective we discussed some of the many security capabilities built into information systems. Specific to this objective, we discussed hardware and firmware system security that makes use of Trusted Platform Modules, hardware security modules, self-encrypting drives, and bus encryption. We also discussed characteristics of secure processing that rely on this hardware and firmware. We explored the concept of a trusted execution environment, which provides an enclosed, trusted environment for software to execute. We also mentioned processor security extensions, additional code built into a CPU that allows for memory reservation and encryption for sensitive processes. Finally, we talked about atomic execution, which is an approach used to secure or lock a section of code against asynchronous or timing attacks that could interfere with sequencing and task execution timing.

3.4 QUESTIONS

1. Which of the following security capabilities is able to encrypt media and can be moved from system to system, but does not rely on the internal device cryptographic mechanisms to manage device encryption?
   - **A.** Hardware security module
   - **B.** Self-encrypting drive
   - **C.** Bus encryption
   - **D.** Trusted Platform Module
2. Which of the following secure processing capabilities is used to protect against time-of-check to time-of-use attacks?
   - **A.** Trusted execution environment
   - **B.** Trusted Platform Module
   - **C.** Atomic execution
   - **D.** Processor security extensions

3.4 ANSWERS

1. **B** A self-encrypting drive has its own cryptographic hardware mechanisms built into the drive electronics, so it does not rely on other device security mechanisms, such as TPMs, HSMs, or bus encryption to manage its encryption capabilities. Additionally, self-encrypting drives can be moved from device to device.
2. **C** Atomic execution is an approach used in secure software construction that locks or isolates specific code segments and prevents outside processes or applications from interrupting their processing and taking advantage of sequencing and timing issues when processes or tasks are executed.



Objective 3.5 Assess and mitigate the vulnerabilities of security architectures, designs, and solution elements

In this objective we will take a look at various security architectures and designs, examining what they are and how they may fit into an overall organizational security picture. We will also discuss some of the vulnerabilities that affect each of these architectures.

Vulnerabilities of Security Architectures, Designs, and Solutions

There are a multitude of ways that infrastructures can be designed and architected, and each way has its own advantages and disadvantages as well as unique security issues. In this objective we will discuss several of the key architectures that you will see both on the exam and in the real world. We will also talk about some of the vulnerabilities that plague each of these solutions. Most of the vulnerabilities associated with all of these solutions involve recurring themes: weak authentication, lack of strong encryption, lack of restrictive permissions, and, in the case of specialized devices, the inability to securely configure them while still connecting them to sensitive networks or even the larger Internet.

Client-Based Systems

A client-based system is one of the simplest computer architectures. It usually does not depend on any external devices or processing power; all physical hardware and necessary software are self-contained within the system. Client-based systems use applications that are executed entirely on a single device. The device may or may not have or need any type of network connectivity to other devices.

Although client-based systems are not the norm these days, you will still see them in specific implementations. In some cases, they are simply older or legacy devices and implementations, and in many other cases they are devices that have very specialized uses and are designed intentionally to be independent and self-contained. Regardless of where you see them or why they exist, client-based systems still may require patches and updates and, due to limited storage, may require connections to external storage devices. Occasional external processing could involve connections to other applications or network-enabled devices. However, all the core processing is still performed on the client-based system.

Vulnerabilities associated with client-based systems are the same as those that are magnified to a greater extent on other types of architectures. Client-based systems often suffer from weak authentication because the designers do not feel there is a risk in using simple (or sometimes nonexistent) authentication, based on the assumption that the systems likely won't connect to anything else. There also may be very weak encryption mechanisms, if any at all, on the device. Again, the assumption is that the device will not connect to any other devices, so data does not need to be transmitted in an encrypted form. However, this assumption does not consider data that should be encrypted while in storage on the device. Most security mechanisms on the client-based system are a single point of failure; that is, if they fail, there is no backup or redundancy.

Server-Based Systems

Server-based system architectures extend client-based system architectures (where everything occurs on a single device or application) by connecting clients to the network, enabling them to communicate with other systems and access their resources. Server-based systems, commonly called a client/server or two-tiered architecture, consist of a client (either a device

or an application) that connects to a server component, sometimes on the same system or another system on the network, to access data or services. Most of the processing occurs on the server component, which then passes data to the client.

Vulnerabilities inherent to server-based systems include operating system and application vulnerabilities, weak authentication, insufficient encryption mechanisms, and, in the case of network client/server components, nonsecure communications.

Distributed Systems

Distributed systems are those that contain multiple components, such as software or data residing on multiple systems across the network or even the Internet. These are called *n-tier* computing architectures, since there are multiple physical or software components connected together in some fashion, usually through network connections. A single tier is usually one self-contained system, and multiple tiers means that applications and devices connect to and rely on each other to provide data or services. These architectures can range from a simple two-tiered client/server model, such as one where a client's web browser connects to a web server to retrieve information, to more complex architectures with multiple tiers and many different components, such as application or database servers.

What characterizes distributed systems the most is that processing is shared between hosts. Some hosts in an n-tiered system provide most of the processing, some hosts provide storage capabilities, and still others provide services such as security, data transformation, and so on.

There are many different security concerns with n-tiered architectures. One is the flow of data between components. From strictly a functional perspective, data latency, integrity, and reliability are concerns since connections can be interrupted or degraded, which then affects the availability goal of security. However, other security concerns include nonsecure communications, weak authentication mechanisms, and lack of adequate encryption. Two other key considerations in n-tiered architectures are application vulnerabilities and vulnerabilities associated with the operating system on any of the components.

Database Systems

Databases come in many implementations and are often targets for malicious entities. A database can be a simple spreadsheet or desktop database that contains personal financial information, or it can be a large-scale, multitiered big data implementation. There are several models for database construction, and some of these lend themselves to security better than others.

Relational databases, which make up the large majority of database design, are based on *tables* of information, with rows and columns (also called *records* and *fields*, respectively). Database *rows* contain data about a specific subject, and *columns* contain specific information elements common to all subjects in the table. The *key* is one or more fields in the table used to uniquely identify a record and establish a relationship with a related record in a different table. The *primary key* is a unique combination of fields that identifies a record based upon nonrepeating data. A *foreign key* is used to establish a relationship with a related record

in a different table or even another database. *Indexes* are keys that are used to facilitate data searches within the database.

In addition to relational databases (the most common type), there are other database architectures that exist and are used in different circumstances. Some of these are simple databases that use text files formatted in simple comma- or tab-separated fields and use only one table (called a *flat-file* database). Another type, called *NoSQL*, is used to aggregate databases and data sources that are disparate in nature, such as a structured database combined with unstructured (unformatted) data. NoSQL databases are used in "big data" applications that must glean usable information from multiple data sources that have no format, structure, or even data type in common. A third type of database architecture is the *hierarchical* database, which is an older type still used in some applications. Hierarchical databases use a hierarchical structure where a main data element has several nested hierarchies of information under it. Note that there are many more types of database architectures and variations of the ones we just mentioned.

Regardless of architecture, common vulnerabilities of databases include poor design that allows end users to view information they should not have access to, inadequate interfaces, lack of proper permissions, and vulnerabilities in the database management system itself. To combat these vulnerabilities, implement restrictive views and constrained interfaces to limit the data elements that authorized users are able to access, as well as database encryption and vulnerability/patch management for database management systems.

Cryptographic Systems

Cryptographic systems have two primary types of vulnerabilities: those that are inherent to the cryptographic algorithm or key, and those that affect how a cryptosystem is implemented. Between the two, weak algorithms and keys are more common, but there can be issues with how cryptographic systems are designed and implemented. We will go more into depth on the weaknesses of cryptographic systems in Objective 3.6.

Industrial Control Systems

Industrial control systems (ICSs) are a unique set of information technologies designed to control physical devices in mechanical and physical processes. ICSs are different from traditional IT systems in that some of their devices, programming languages, and protocols are legacy or proprietary and were meant to stand alone and not connect to traditional networks. A *distributed control system (DCS)* is a network of these devices that are interconnected but may be part of a closed network; typically a DCS consists of devices that are in close proximity to each other, so they do not require long-haul or wide-area connections. A DCS may be the type of control device network encountered in a manufacturing or power facility. Whereas a DCS is limited in range and has devices that are very close to each other, a *supervisory control and data acquisition (SCADA)* system covers a large geographic area and consists of multiple physical devices using a mixture of traditional IT and ICS.



NOTE Both DCS and SCADA systems, as well as some IoT and embedded systems, all fall under the larger category of ICS.

Concepts and technologies related to ICSs that you should be aware of for the exam include

- **Programmable logic controllers (PLCs)** These are devices that control electromechanical processes.
- **Human–machine interfaces (HMIs)** These are monitoring systems that allow operators to monitor and interact with ICS systems.
- **Data historian** This is the system that serves as a data repository for ICS devices and commonly stores sensor data, alerts, command history, and so on.

Because ICS devices tend to be relatively old or proprietary in nature, they often do not have authentication or encryption mechanisms built-in, and they may have weak security controls, if any. However, modern technologies and networks now connect to some of these older or legacy devices, which poses issues regarding security compatibility and the many vulnerabilities involving data loss and the ability to connect into otherwise secure networks through these legacy devices. Operational technology (OT) is the grouping together of traditional IT and ICS devices into a system and requires special attention in securing those networks.

Internet of Things

Closely related to ICS devices are the more modern, commercialized, and sometimes consumer-driven versions of those devices, popularly referred to as the Internet of Things (IoT). IoT is a more modern approach to embedding intelligent systems that can interact with and connect to other systems, including the worldwide Internet. A wide range of devices—from smart refrigerators, doorbells, and televisions to medical devices, wearables (e.g., smart watches), games, and automobiles—are considered part of the Internet of Things, as long as they have the ability to connect to other systems via wireless or wired connections. IoT devices use standardized communications protocols, such as TCP/IP, and have very specialized but sometimes limited processing power, memory, and data storage. Inadequate security when connected to the Internet is their primary weakness, as many of these IoT devices do not have advanced authentication, encryption, or other security mechanisms, which makes them an easy entryway into other traditional systems that may house sensitive data.

Embedded Systems

Embedded systems are integrated computers with all of their components self-contained within the system. They are typically designed for specific uses and may only have limited processing or storage capabilities. Examples of embedded systems are those that control engine functions in a modern automobile or those that control aircraft systems. Embedded systems

are similar to ICS/SCADA and IoT systems in that they are special-function devices that are often connected to the Internet, sometimes without consideration for security mechanisms. Many embedded systems are proprietary and do not have robust, built-in security mechanisms such as strong authentication or encryption capabilities. Additionally, the software in embedded systems is often embedded into a computer chip and may not be easily updatable or patched as vulnerabilities are discovered for the system.

Cloud-Based Systems

Cloud computing is a set of relatively new technologies that facilitate the use of shared remote computing resources. A cloud service provider (CSP) offers services to clients that subscribe to the services. Normally, the cloud service provider owns the physical hardware, and sometimes the infrastructure and software, that is used by the clients to access the services. A client connects to the provider's infrastructure remotely and uses the resources as if the client were connected on the client's premises. Cloud computing is usually implemented as large data centers that run multiple virtual machines on robust hardware that may or may not be dedicated to the client.

Organizations subscribe to cloud services for many reasons, which include cost savings by reducing the necessity to build their own infrastructures, buy their own equipment, or hire their own network and security personnel. The cloud service provider takes care of all these things, to one degree or another, based upon what the client's subscription offers.

There are three primary models for cloud computing subscription services:

- **Software as a Service (SaaS)** The client subscribes to applications offered by the CSP.
- **Platform as a Service (PaaS)** Virtualized computers run on the CSP's infrastructure and are provisioned for the use of the client.
- **Infrastructure as a Service (IaaS)** The CSP offers networking infrastructure and virtual hosts that clients can provision and use as they see fit.

Other types of cloud services that CSPs offer include Security as a Service (SECaaS), Database as a Service (DBaas), Identity as a Service (IDaas), and many others.

In addition to the three primary subscription models, there are also four primary deployment models for cloud infrastructures:

- **Private cloud** The client organization owns the infrastructure, or subscribes to a dedicated portion of infrastructure located in a cloud service provider's data center.
- **Public cloud** The client organization uses public cloud infrastructure providers and subscribes to specific services only; infrastructure services are not dedicated solely to the client but are shared among different clients.
- **Community cloud** Cloud infrastructure is shared between like organizations for the purposes of collaboration and information exchange.
- **Hybrid cloud** This is a combination of any or all of the other deployment models, and is likely the most common you will see.

While a cloud service provider uses a shared responsibility model that offsets some of the burden of security protection from the client, there are still vulnerabilities in cloud-based models and deployments. Some of the considerations for the shared security model include

- **Delineated responsibilities** Responsibilities between the cloud service provider and the client must be delineated and accountability must be maintained.
- **Administrative access to devices and infrastructure** The client may or may not have access to all of the devices and infrastructure they use and may have limited ability to configure or secure those devices.
- **Auditing** Both the client and the CSP must share responsibility for auditing actions on the part of client personnel and the CSP.
- **Configuration and patching** Usually the CSP is responsible for host, software, and infrastructure configuration and patching, but this must be detailed in the agreement.
- **Data ownership and retention** Data owned by the client, particularly sensitive data such as healthcare information and personal data, may be accessed or processed by the CSP's personnel, but accountability for its protection remains with the client.
- **Legal liability** CSP agreements must address issues such as breaches in terms of legal liability, responsibility, and accountability.
- **Data segmentation** The client must determine if the data will be comingled with or segmented from other organizations' data when it is processed and stored on the CSP's equipment.

Virtualized Systems

Virtualized systems use software to emulate hardware resources; they exist in simulated environments created by software. The most popular example of a virtualized system is a virtual operating system that is created and managed by a hypervisor. A *hypervisor* is responsible for creating the simulated environment and managing hardware resources for the virtualized system. It acts as a layer between the virtual system and the higher-level operating system and physical hardware. A hypervisor comes in two flavors:

- **Type I** Also called a *bare-metal hypervisor*, this is a minimal, special-purpose operating system that runs directly on hardware and supports virtual machines installed on top of it.
- **Type II** This is implemented as a software application that runs on top of an existing operating system, such as Windows or Linux, and acts as an intermediary between the operating system and the virtualized system. The virtualized OS is sometimes referred to as a *guest*, and the physical system it runs on is referred to as a *host*.

Containerization

A virtualized system can be an entire computer, including its operating system, applications, and user data. However, sometimes a full virtualized system is not needed. Virtualized systems can be scaled to a much smaller level than guest operating systems. A smaller simulated environment, called a *container*, can be created that simply runs an application in its entirety so there's no need for a full virtual operating system. The container interacts with the hypervisor to get all the necessary resources, such as CPU processing time, RAM, and so on. This allows for minimal resource use and eliminates the necessity to build an entire virtual computer. Popular containerization software includes Kubernetes and Docker.

Microservices

Microservices are a minimalized form of containerization that doesn't require building a large application. Application functionality and services that an application might otherwise provide are divided up into much smaller components, called *microservices*. Microservices run in a containerized environment and are essentially small, decentralized individual services designed and built to support business capabilities. Microservices are also independent; they tend to be loosely coupled, meaning they do not have a lot of required dependencies between individual services. Microservices can be a quick and efficient way to rapidly develop, test, and provision a variety of functions and services.

Serverless

An even more minimalized form of virtualization are serverless functions. In a serverless implementation, services such as compute, storage, messaging, and so on, along with their configuration parameters, are deployed as microservices. They are called *serverless functions* because a dedicated hosting service or server is not required. Serverless architectures work at the individual function level.

High-Performance Computing Systems

High-performance computing (HPC) systems are far more powerful than traditional general-purpose computing systems, and they are designed to solve large problems, such as in mathematics, economics, and advanced physics. In addition to having extremely high-end hardware in every step of the processing chain, HPC systems also aggregate computing power by using multiple processors (sometimes hundreds of thousands), faster buses, and specialized operating systems designed for speed and accuracy.

Edge Computing Systems

In the age of real-time financial transactions, video streaming, high-end gaming, and the need to have massive processing power and storage with lightning-fast network speeds readily available, edge computing has become a more important piece of distributed computing.

Edge computing is the natural evolution of *content distribution networks* (*CDNs*; also sometimes referred to as *content delivery networks*), which were originally invented to deliver web content (think gaming and video streaming services) to users. As large data centers hosting these services became more powerful, and then distributed throughout large regions, slow, cross-country wide area network (WAN) links often became the bottleneck. It became a necessity to bring speed and efficiency closer to the user.



EXAM TIP CDNs provide redundance and more reliable delivery of services by locating content across several data centers. Edge computing is designed to bring that content geographically closer to the user to overcome slow WAN links. Both, however, use similar methods and are almost part of the same infrastructures.

Rather than a user connecting through a network of potentially slow links across an entire country or even overseas, edge computing allows for intermediate distribution points for data and services to be established physically and logically closer to the user to help reduce the dependency on overtaxed long-haul connections. This way, a user doesn't necessarily have to maintain a constant connection to a centralized data center for video streaming; the content can also be replicated to edge computing points so that the user can simply access those points without having to contend with slower links over large distances. Edge computing also uses the fastest available equipment and links to help minimize latency.

REVIEW

Objective 3.5: Assess and mitigate the vulnerabilities of security architectures, designs, and solution elements In this objective we discussed different types of computing architectures and designs, as well as their vulnerabilities. Client-based systems do not depend on any external devices or processing power and are not necessarily connected to a network. Server-based systems have client and server-based components to them. N-tier architectures are distributed systems that use multiple components. We looked at the basics of database systems, including relational and hierarchical systems. Nontraditional IT systems include industrial control systems, SCADA systems, and Internet of Things devices which may or may not have secure authentication or encryption mechanisms built in. We also reviewed cloud subscription services and cloud deployment architectures, to include Software as a Service, Infrastructure as a Service, and Platform as a Service, as well as public, private, community, and hybrid cloud models. Virtualized systems are made up of various components including hypervisors, virtualized guests, physical hosts, containers, microservices, and serverless architectures. These components emulate not only for operating systems, but also applications and lower-level constructs to minimize hardware requirements and provide specific functions. We also briefly discussed embedded systems which are essentially systems embedded into computer chips, as well as computing

platforms on the opposite end of the scale which deal with massive amounts of processing power, or high-performance computing. Finally, we briefly discussed edge computing systems, which deliver services closer to the user to help eliminate latency issues over wide area networks.

3.5 QUESTIONS

1. A large system in a company has many components, including an application server, a web server, and a backend database server. Clients access the system through a web browser. Which of the following best describes this type of architecture?
   - **A.** N-tier architecture
   - **B.** Client/server architecture
   - **C.** Client-based system
   - **D.** Serverless architecture
2. Your company needs to provide some functionality for users that perform only very minimal services but will connect to another, larger line of business applications. You don't need to program another large enterprise-level application. Which of the following is the best solution that will fit your needs?
   - **A.** Microservices
   - **B.** Virtualized operating systems
   - **C.** Embedded systems
   - **D.** Industrial control systems

3.5 ANSWERS

1. **A** An n-tier architecture is characterized by multiple distributed components. In this case, the n-tier architecture is composed of an application server, web server, and database server, along with its client-based web browsing connection.
2. **A** Microservices provide low-level functionality that can be accessed by other applications, without the need to build an entire enterprise-level application.



Objective 3.6 Select and determine cryptographic solutions

Objectives 3.6 and 3.7 cover cryptography. You should understand the basic terms and concepts associated with cryptography, but you don't have to understand the math behind it for the exam. We will go over basic terms and concepts in this objective, and in Objective 3.7 we will discuss some of the attacks that can be perpetrated on cryptography.

In this objective we will look at various aspects of cryptology and cryptography, including the cryptographic life cycle, which explains how keys and algorithms are selected and managed, and cryptographic methods that are used, such as symmetric, asymmetric, and quantum cryptography. We'll also look at the application of cryptography, in the form of the public key infrastructure (PKI).

Cryptography

Cryptography is the science of storing and transmitting information that can only be read or accessed by specific entities through the use of mathematical algorithms and keys. While cryptography is the most popular term we use in association with the science, *cryptology* is the overarching term that applies to both cryptography and cryptanalysis. *Cryptanalysis* refers to analyzing and reversing cryptographic processes in order to break encryption. *Encryption* is the process of converting *plaintext* information, which can be read by humans and computers easily, into what is called *ciphertext*, which cannot be read or accessed by humans or machines without the proper cryptographic mechanisms in place. *Decryption* reverses that process and allows authorized entities to view information hidden by encryption. Cryptology supports the confidentiality and integrity goals of security, and also helps to ensure that supporting tenets, such as authentication, nonrepudiation, and accountability, are met.

Most cryptography in use today involves the use of *ciphers*, which convert individual characters (or the binary bits that make up the characters) into ciphertext. By contrast, the term *code* refers to using symbols, rather than characters or numbers, to represent entire words or phrases. Earlier ciphers used simple *transposition* to rearrange the characters in a message, so that they were merely scrambled among themselves, or *substitution*, in which the cipher replaced characters with different characters in the alphabet.



NOTE The term *cipher* can also be interchanged with the term *algorithm*, which refers to the mathematical rules that a given encryption/decryption process must follow. We'll discuss algorithms in the next two sections as well.

Cryptographic Life Cycle

There are many different components to cryptography, each of which is controlled according to a carefully planned life cycle. The *cryptographic life cycle* is the continuing process of identifying cryptography requirements, choosing the algorithms that meet your needs, managing keys, and implementing the cryptosystems that will facilitate the entire encryption/decryption process. Regardless of the method used for cryptography, there are three important components: algorithms, keys, and cryptosystems.

Algorithms

Algorithms are complex mathematical formulas or functions that facilitate the cryptographic processes of encryption and decryption. They dictate how plaintext is manipulated to produce ciphertext. Algorithms are normally standardized and publicly available for examination. If it were just a matter of putting plaintext through an algorithm and producing ciphertext, this process could be repeated over and over to produce the same resulting ciphertext. However, this predictability eventually would be a vulnerability in that, given enough samples of plaintext and ciphertext, someone may figure out how the algorithm works. That's why another piece of the puzzle, the key, is introduced to add variability and unpredictability to the encryption process.

Keys

The *key*, sometimes called a *cryptovariable*, is introduced into the encryption process along with the algorithm to add to the complexity of encryption and decryption. Keys are similar to passwords in that they must be changed often and are usually known only to the authorized entities that have access and authorization to encrypt and decrypt information. In 1883, Auguste Kerckhoffs published a paper stating that only the key should be kept secret in a cryptographic system. Known as *Kerckhoffs' principle*, it further states that the algorithm in a cryptographic system should be publicly known so that its vulnerabilities can be discovered and mitigated.

There have been some exceptions to this principle; of note are the U.S. government's Clipper Chip and Skipjack algorithms, whose inner workings were kept secret under the theory that if no one outside the government circles knew how they worked, then no one would be able to discover vulnerabilities. Industry and professional security communities disagree with this approach, as it tends to follow the faulty principle of *security through obscurity*, meaning that simply not being aware of a security control makes it stronger. As such, most commonly used algorithms today are open for inspection and review.

Keys should be created to be strong, meaning that the following guidance should be considered when creating keys:

- Longer key lengths contribute to stronger keys.
- A larger key space (number of available character sets and possibilities for each in a key) creates stronger keys.
- Keys that are more random and not based on common terms, such as dictionary words, are stronger.
- Keys should be protected for confidentiality the same as passwords—not written down in publicly accessible places or shared.
- Keys should be changed frequently to mitigate against attacks that could compromise them.
- No two keys should generate the same ciphertext from the same plaintext, which is called *clustering*.

Cryptosystems

A cryptosystem is the device or mechanism that implements cryptography. It is responsible for the encryption and decryption processes. A cryptosystem does not have to be a piece of hardware; it can be software, such as an application that generates keys or encrypts e-mail. Cryptosystems are built in a secure fashion, but also use complex keys and algorithms. The strength of a cryptosystem is called its *work function*. The higher the work function, the stronger the cryptographic system. As we will see in Objective 3.7, attacks against cryptosystems and any weaknesses in their implementation are commonplace.

Cryptographic Methods

Using the common cryptographic components of algorithms, keys, and cryptosystems, cryptography is implemented using many different methods and means. Cryptographic methods are designed to increase strength and resiliency of cryptographic processes while at the same time reducing their complexity when possible.

Two basic operations are used to convert plaintext into ciphertext:

- **Confusion** The relationship between any plaintext and the key is so complex that an attacker cannot encrypt plaintext into ciphertext and figure out the key from those relationships.
- **Diffusion** A very small change in the plaintext, even a single character in a sentence, causes much larger changes to ripple through the resulting ciphertext. So, changing a single character in the plaintext does not necessarily result in just changing a single character in the ciphertext; the resulting change would likely change a large part of the resulting ciphertext.

Another cryptographic process you should be aware of involves the use of Boolean math, which is binary. Cryptography uses logical operators, such as AND, NAND, OR, NOR, and XOR, to change binary bits of plaintext into ciphertext. The exact combination of logical operations used, and their order, depends largely on the algorithm in use. These processes also use one-way functions. A *one-way function* is a mathematical operation that cannot be easily reversed, so plaintext that processes through one of these functions into ciphertext cannot be reversed to its original state from ciphertext using the same function, without knowing the algorithm and key. Some algorithms also add random numbers or seeds (called a *nonce*) to make the encryption process more complex, random, and difficult to reverse engineer.

Various cryptographic methods exist, but for purposes of preparing for the CISSP exam, you should be familiar with four in particular: symmetric encryption, asymmetric encryption, quantum cryptography, and elliptic curve cryptography.

Symmetric Encryption

Symmetric encryption (also called *secret key* or *session key* cryptography) uses one key for its encryption and decryption operations. The key is selected by all parties to the communications process, and everyone has the same key. The key can be generated on-the-fly by the

cryptographic system or application, such as when a person accesses a secure website and negotiates a secure connection with the remote server. The main characteristics of symmetric encryption include the following:

- It is not easily scalable, since complete confidential communications between all parties involved requires many more keys that have to be managed.
- It is relatively fast and suitable for encrypting bulk data.
- Key exchange is problematic, since symmetric keys must be given to all parties securely to prevent unauthorized entities from getting access to the key.

One of the main vulnerabilities in symmetric encryption is that it is not scalable, since the number of keys required between multiple parties increases as the size of the group requiring encryption among themselves increases. For example, if two people want to use symmetric encryption to send encrypted messages between each other, they need only one symmetric key. However, if three people are involved, and they each need to send encrypted messages between each other that only the other person can decrypt, each person needs two symmetric keys to connect to the other two. For larger groups, this number can become unwieldy; the formula for determining the number of keys needed between multiple parties is $N(N - 1)/2$, where N is the number of parties involved.

As an example, if 100 people need to be able to exchange secure messages with each other, and all 100 people are authorized to read all messages, then only one symmetric key is needed. However, if certain messages have to be encrypted and only decrypted by certain persons within that group, then more keys are needed. In this example, to ensure that each individual can exchange secure messages with every other single individual only, according to the formula $100(100 - 1)/2$, then 4,950 individual keys would be required. As you will see in the next section, asymmetric encryption methods can solve the scalability problem.

Symmetric encryption methods use either of two primary types of cipher: block or stream. A *block* cipher encrypts data in huge chunks, called blocks. Block sizes are measured in bits. Common block sizes are 64-bit, 128-bit, and 256-bit blocks. Typical block ciphers include Blowfish, Twofish, DES, and the Advanced Encryption Standard, or AES, which are detailed in Table 3.6-1.

Stream ciphers, on the other hand, encrypt data one bit at a time. Stream ciphers are much faster than block ciphers. The most common example of a widely used stream cipher is RC4. Initialization vectors (IVs) are random seed values that are used to begin the encryption process with a stream cipher.

Asymmetric Encryption

Asymmetric encryption, unlike its symmetric counterpart, uses two keys. One of these keys is arbitrarily referred to as the public key, and the other is the private key. The keys are mathematically related but not identical. Having access to or knowledge of one key does not allow

TABLE 3.6-1 Common Symmetric Algorithms

Algorithm	Characteristics	Notes
DES	Uses 64-bit blocks, 16 rounds of encryption; key size is 56 bits	Lucifer algorithm selected as the U.S. Data Encryption Standard
3DES	Uses 64-bit blocks; 16 rounds of encryption; 168-bit key (three 56-bit keys)	Repeats the DES process three times
AES	Uses 128-bit blocks; 10, 12, and 14 rounds of encryption (based on key size); 128-, 192-, and 256-bit keys	Uses the Rijndael algorithm selected as the U.S. Advanced Encryption Standard
Blowfish	Uses 64-bit blocks; 16 rounds of encryption; 32- to 448-bit keys	Public domain algorithm
Twofish	Uses 12-bit blocks; 16 rounds of encryption; 128-, 192-, and 256-bit keys	Public domain algorithm
RC4	Streaming cipher, uses one round of encryption	Used in WEP, SSL, and TLS (largely deprecated in current technologies)

someone to derive the other key in the pair. An important thing to understand about asymmetric encryption is that whatever data one key encrypts, only the other key can decrypt, and vice versa. You cannot decrypt ciphertext using the same key that was used to encrypt it. This is an important distinction to make since it is the foundation of asymmetric encryption, also referred to as public key cryptography.

Other key characteristics of asymmetric encryption include the following:

- It uses mathematically complex algorithms to generate keys, such as prime factorization of extremely large numbers.
- It is much slower than symmetric encryption.
- It is very inefficient at encrypting large amounts of data.
- It is very scalable, since even a large group of people exchanging encrypted messages only each need their own key pair, not multiple keys from other people.

Asymmetric encryption allows the user to retain one key, referred to as the user's *private key*, and give the *public key* out to anyone. Since anything encrypted with the public key can only be decrypted with the user's private key, this ensures confidentiality. Conversely, anything encrypted with a user's private key can only be decrypted with the public key. While this definitely does not guarantee confidentiality, it does assist in ensuring authentication, since only the user with the private key could have encrypted data that anyone else with the public key can decrypt. Common asymmetric algorithms include those listed in Table 3.6-2.

TABLE 3.6-2 Common Asymmetric Algorithms

Algorithm	Characteristics	Notes
Rivest, Shamir, Adleman (RSA)	Common key sizes from 1,024 to 4,096 bits	De facto algorithm used in generating public/private key pairs
Diffie-Hellman (DHE)	Basis of modern-day public key cryptography	Series of key exchange protocols
ECC	Based on the algebraic structure of elliptic curves over finite fields	Requires far less computing power than other algorithms; suitable for mobile devices
El Gamal	Based on the difficulty of computing certain problems in discrete logarithms	Used for general encryption and digital signatures; based partially on Diffie-Hellman

Quantum Cryptography

Quantum cryptography is a cutting-edge technology that uses quantum mechanics to provide for cryptographic functions that are essentially impossible to eavesdrop on or reverse engineer. Although much of quantum cryptography is still theoretical and years away from practical implementation, *quantum key distribution (QKD)* is one aspect of quantum cryptography that is becoming more useful in the near term. QKD can help with secure key distribution issues associated with symmetric cryptography. It uses orientation of polarized photons, assigned to binary values, to pass keys from one entity to another. Note that observing or attempting to measure the orientation of these photons actually changes them and disrupts communication, rendering the process moot. This is why it would be very easy to detect any unauthorized eavesdropping or modification of the communication stream.

Elliptic Curve Cryptography

Elliptic curve cryptography (ECC) is an example of an asymmetric algorithm and uses points along an elliptic curve to generate public and private keys. The keys are mathematically related, as are all asymmetric keys, but cannot be used to derive each other. ECC is used in many different functionalities, including digital signatures (discussed in an upcoming section), secure key exchange and distribution, and, of course, encryption. The main characteristics of ECC you need to remember are that it is highly efficient, much more so than other asymmetric algorithms, is far less resource intensive, and requires much less processing capability. This makes ECC highly suitable for mobile devices that are limited in processing power, storage, electrical power, and bandwidth.

Integrity

Along with confidentiality and availability, integrity is one of the three primary goals of security, as described in Objective 1.2. Integrity ensures that data and systems are not tampered with, and that no unauthorized modifications are made to them. Integrity can be assured

in a number of ways, one of which is via cryptography, which establishes that data has not been altered by computing a one-way hash (also called a *message digest*) of the piece of data. For example, a message that is hashed using a one-way hashing function produces a unique fingerprint, or message digest. Later, after the message has been transmitted, received, or stored, if the integrity of the message is questioned or required, the same hashing algorithm is applied to the message. If the message has been unaltered, the hash should be identical. If it has been changed in any way, either intentionally or unintentionally, the resulting hash will be different.

Hashing can be used in a number of circumstances, including data transmission and receipt, e-mail, password storage, and so on. Remember that hashing is not the same thing as encryption; the goal of hashing is not to encrypt something that must be decrypted. Hashes are generated using a one-way function that cannot be reversed and then the hashes are compared, not decrypted.

Hash values should be unique to a piece of data and not duplicated by any other different piece of data using the same algorithm. If this occurs, it is called a *collision*, and represents a vulnerability in the hashing algorithm. Popular hashing algorithms include MD5 and the Secure Hash Algorithm (SHA) family of algorithms.



EXAM TIP Although it is a cryptographic process, hashing is not the same thing as encryption. Encrypted text can be decrypted; that is, reversed back to its plaintext state. Hashed text is not decrypted, the text is merely hashed again so the hashes can be compared to verify integrity.

Hybrid Cryptography

Hybrid cryptography, as you might guess from its name, is the combination of multiple methods of cryptography, primarily symmetric and asymmetric. On one hand, symmetric encryption is not easily scalable but is quick and well suited for encrypting bulk amounts of information; it's not suitable for use among large groups of people. On the other hand, asymmetric encryption only uses two keys per person, and the public key can be given to anyone, which makes it very scalable. But it is slower and not suitable for bulk data encryption, which is more appropriate for symmetric encryption. It's easy to see that each of these algorithms makes up for the disadvantages of the other, so it makes sense to use them together.

In order to use asymmetric and symmetric encryption together two people who wanted to share encrypted data between them could do the following:

1. The first user (the sender) would give their public key to the second user (the recipient).
2. The recipient would generate a symmetric key, and then encrypt that symmetric key with the sender's public key.
3. The sender would then decrypt the encrypted key using their own private key.

Once the sender has the key, then neither party needs to worry about public and private key pairs; they can simply use the symmetric key to exchange large amounts of data quickly. In this particular example, the symmetric key is called a *session key*, since it is used only for that particular session of exchanging encrypted data.

The asymmetric encryption method allows for secure key exchange, which also happens in the practical world when a user establishes an HTTPS connection to a secure server using a web browser. In this example, one party (usually the client side) creates a session key for both parties to use. Since it's not secure to send that key to the second party (in this case, the server) in an unencrypted form, the sender can use the recipient's public key to encrypt the session key. The recipient then uses their private key to decrypt the message, allowing them to access the session key. Now that both parties have the session key, they can use it to exchange large amounts of data in a fast and efficient manner.

Digital Certificates

Digital certificates are electronic files that contain public and/or private keys. The digital certificate file is a way of securely distributing public and private keys and serves to prove they are solidly connected to an individual entity, such as a person or organization. The certificates are generated through a server that issues them, called a *certificate authority* server. Digital certificates can be installed in any number of software applications or even hardware devices. Once installed, they reside in the system's secure certificate store. This enables them to be used automatically and transparently when users need them to encrypt e-mail, connect to a secure financial website, or encrypt a file before transmitting it.

Digital certificates contain attributes such as the user's name, organization, the certificate issuance date, its expiration date, and its purpose. Digital certificates can be used for a variety of purposes, and a digital certificate can be issued for a single purpose or multiple purposes, depending upon the desires of the issuing organization. For example, a digital certificate could be used to encrypt e-mail and files, provide for identity authentication, and so on, or a digital certificate could be generated only for a specific purpose and used only by a software development company to digitally sign software. Most digital certificates use the popular X.509 standard, which dictates their format and structure. Common file formats for digital certificates include DER, PEM, and the Public Key Cryptography Standards (PKCS) #7 and #12.



NOTE PKCS was developed by RSA Security and is a proprietary set of standards.

Public Key Infrastructure

Public key infrastructure, or PKI, is the name given to the set of technologies, processes, and procedures, along with associated infrastructure, that implements hybrid cryptography (using both asymmetric and symmetric methods and algorithms) to provide a trusted infrastructure

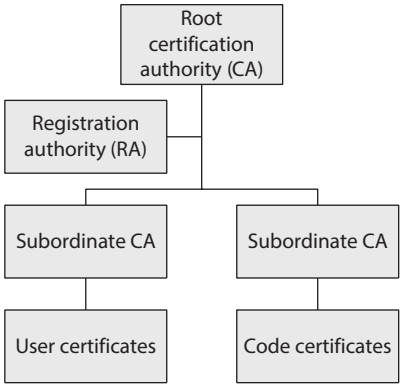


FIGURE 3.6-1 A typical PKI

for encryption, file integrity, and authentication. PKI follows a formalized hierarchical structure, as shown in Figure 3.6-1.

The typical PKI consists of a trusted entity, called a *certificate authority (CA)*, that provides the certificates and keys. This may be a commercial company that specializes in this business (such as Verisign, Thawte, Entrust, etc.) or even a department within your own company that is charged with issuing keys and certificates. Trusted entities can also be third parties, just as long as they can validate the identities of subjects. Trust in a CA is based on the assumption that it verifies the identity and validity of entities to which it issues keys and certificates (called *subjects*).

Generally, if a certificate is to be used or trusted only within your organization, it should come from an internal CA; if you need organizations external to yours to trust the certificate (e.g., customers, business partners, suppliers, etc.), then you should use an external or third-party CA.



NOTE The term *certificate authority* can refer to either the trusted entity itself or to the server that actually generates the keys and certificates, depending on the context.

In larger organizations, the workload of validating identities is often offloaded to a different department or organization; this is known as the *registration authority (RA)*. The RA collects proof of identity, such as an individual's driver's license, passport, birth certificate, and so on, and validates that the user is actually who they say they are. The RA then passes this information on to the CA, which performs the actual key generation and issues the digital certificates.

In organizations where the *root* (top-level) CA performs all the key-generation and certificate-issuance functions, it might be prudent to add *subordinate* (sometimes referred to as *intermediate*) CAs, which are additional servers tasked with issuing specific types of certificates to offload the root CA's work. The subordinate servers could issue specific certificates, such

as user certificates or code signing certificates, or they could simply share the load of issuing all the same certificates as the root CA. Implementing subordinate CAs is also a good security practice. By having subordinate CAs issue all certificates, the root CA can be taken offline to protect it, since an attack on that server would compromise the entire trust chain of the certificates issued by the organization.

While PKI is the most common implementation of asymmetric cryptography, it is not the only one. As just described, PKI relies on a hierarchy of trust, which begins with a centralized root CA, and is constructed similar to a tree. Other models implement digital identities, including peer-to-peer models, such as the web-of-trust model used by the Pretty Good Privacy (PGP) encryption program.

Nonrepudiation and Digital Signatures

Nonrepudiation means that someone cannot deny that they took an action, such as sending a message or accessing a resource. Nonrepudiation is vital to accountability and auditing, since accountability is used to ensure that any individual's actions can be traced back to them and they can be held responsible for those actions. Auditing assists in establishing nonrepudiation for a variety of actions, including those relating to accessing resources, exercising privileges, and so on.

Cryptography can assist with nonrepudiation through the use of digital signatures. Digital signatures are completely tied to an individual, so their use guarantees (provided the digital certificate has not been compromised in any way) that an individual cannot deny that they took an action, such as sending an e-mail that has been digitally signed.

Digital signatures support the authenticity process to prove that a message sent by the user is authentic and that the user is the one who sent it. This is a perfect example of hybrid cryptography; the message to be sent is hashed, which provides integrity since the message cannot be altered without changing its hash. The sender's private key encrypts the message hash, and this is sent as part of the transmission to the recipient. When the recipient receives the message (which could also be encrypted in addition to being digitally signed), they can use the sender's public key to decrypt the hash of the message, and then compare the hashes to ensure integrity. Hashing ensures message integrity, and encrypting the hash with the user's private key ensures authenticity, since the message could have come only from the private key holder.

Key Management Practices

Key management consists of the procedures and processes used to manage and protect cryptographic keys throughout their life cycle. Note that keys could be those used by hardware mechanisms or software applications, as well as individual keys issued to users. Management of keys includes setting policies and implementing the key management infrastructure. The practical portion of key management includes

- Verifying user identities
- Issuing keys and certificates
- Secure key storage

- Key escrow (securely maintaining copies of users' keys in the event a user leaves and encrypted data needs to be recovered)
- Key recovery (in the event a user loses a key, or its digital copy becomes corrupt)
- Managing key expiration (renewing or reissuing keys and certificates prior to their expiration)
- Suspending keys (for temporary reasons, such as an investigation or extended vacation)
- Key revocation (permanent action based on circumstances such as key compromise or security policy violations by user)

Among the listed items, one of the most important responsibilities of key management is to ensure that certificates and their associated keys are renewed before they expire. If they are allowed to lapse, it may prove difficult to reuse them, and the organization would have to re-issue new keys and certificates.

Certificate revocation is an important security issue as well. Organizations typically use several means to revoke certificates that are suspected of compromise or misuse. First, an organization publishes a formal *certificate revocation list (CRL)*, which identifies revoked certificates. This process can be manual and produce an actual downloadable list, but most modern organizations use the Online Certificate Status Protocol (OCSP), which automates the process of publishing certificate status, including revocations. The CRL is published periodically or as needed and is also copied electronically to a centralized repository for the organization. This allows for certificates to be checked prior to their use or trust by another organization.

While certificate revocation is normally a permanent action taken in the event of a policy violation or compromise, organizations also have the option to temporarily suspend keys and certificates and then reactivate and reuse them once certain conditions are met. For example, an organization might suspend keys and certificates so that an individual cannot use them during an investigation or during an extended vacation. The organization should always consider carefully whether it needs to revoke a certificate or simply suspend it temporarily, since revoking a certificate effectively renders it permanently void.

Another important consideration is certificate compromise. If this occurs, the organization should revoke the certificate immediately so that it can no longer be used. The details of the revocation should be published to the CRL and sent with the next OCSP update. Additionally, the organization should decrypt all data encrypted with that key pair and issue new keys and certificates to replace the compromised ones.

REVIEW

Objective 3.6: Select and determine cryptographic solutions In this first of two objectives that cover cryptography in depth, we discussed the basic elements of cryptography, including terms and definitions related to cryptography, cryptology, and cryptanalysis, as well as the basic components of cryptography, which include algorithms, keys, and cryptosystems.

We also examined various cryptographic methods, including symmetric cryptography (using a single key for both encryption and decryption), asymmetric cryptography (using a public/private key pair), quantum cryptography, and elliptic curve cryptography.

We briefly discussed the role cryptography plays in ensuring data integrity, in the form of hashing or message digests. Hybrid cryptography is a solution that applies different methods of cryptography, usually symmetric, asymmetric, and hashing, in order to make up for the disadvantages in each of these methods when used individually. With hybrid cryptography, we can ensure confidentiality, integrity, authentication, authenticity, and nonrepudiation. Hybrid cryptography is implemented through the use of digital certificates, which are generated by a public key infrastructure.

Digital certificates are used to ensure authenticity of a message and validate its sender. Public key infrastructure consists of several components, including certificate authorities, registration authorities, and subordinate certificate servers. Key management practices involve managing keys and certificates throughout their life cycle, and include setting policies, issuing keys and certificates, renewing them before they expire, and suspending or revoking them when necessary.

3.6 QUESTIONS

1. You are teaching a CISSP exam preparation class and are explaining the basics of cryptography to the students in the class. Which of the following is the key characteristic of Kerckhoffs' principle?
   - **A.** Both keys and algorithms should be publicly reviewed.
   - **B.** Algorithms should be publicly reviewed, and keys must be kept secret.
   - **C.** Keys should be publicly reviewed, and algorithms must be kept secret.
   - **D.** Neither keys nor algorithms must be publicly reviewed; they should both be kept secret.
2. Evie is a cybersecurity analyst who works at a major research facility. She is reviewing different cryptosystems for the research facility to potentially purchase and she wants to be able to compare them. Which of the following is a measure of a cryptosystem's strength, which would enable her to compare different systems?
   - **A.** Work function
   - **B.** Key space
   - **C.** Key length
   - **D.** Algorithm block size
3. Which of the following types of algorithms uses only a single key, which can both encrypt and decrypt information?
   - **A.** Elliptic curve
   - **B.** Hashing
   - **C.** Asymmetric cryptography
   - **D.** Symmetric cryptography

4. Which of the following algorithms has the disadvantage of not being very effective at encrypting large amounts of data, as it is much slower than other encryption methods?
   - **A.** Quantum cryptography
   - **B.** Hashing
   - **C.** Asymmetric cryptography
   - **D.** Symmetric cryptography

3.6 ANSWERS

1. **B** Kerckhoffs' principle states that algorithms should be open for inspection and review, in order to find and mitigate vulnerabilities, while keys should remain secret.
2. **A** Work function is a measure that indicates the strength of a cryptosystem. It considers several factors, including the variety of algorithms available for the cryptosystem to use, key sizes, key space, and so on.
3. **D** Symmetric cryptographic algorithms only require the use of one key to both encrypt and decrypt information.
4. **C** Asymmetric cryptography does not handle large amounts of data very well when encrypting, and it can be quite slow, as opposed to symmetric encryption, which is much faster and can easily handle bulk data encryption/decryption.



Objective 3.7 Understand methods of cryptanalytic attacks

In this objective we will continue our discussion of cryptography from Objective 3.6 by looking at the many different ways cryptographic systems can be vulnerable to attack, as well as discussing some of those attack methods. This won't make you an expert on cryptographic attacks, but you will become familiar with some of the basic attack methods used against cryptography that are within the scope of the CISSP exam.

Cryptanalytic Attacks

Cryptographic systems are vulnerable to a variety of attacks. *Cryptanalysis* is the process of breaking into cryptosystems. The goal of cryptanalysis, and cryptographic attacks, is to circumvent encryption by discovering the key involved, breaking the algorithm, or otherwise defeating the cryptographic system's implementation such that the system is ineffective.

There are many different attack vectors used in cryptanalysis, but they all target one or more of three primary areas: the key itself, the algorithm used to create the encryption process, and the implementation of the cryptosystem. Secondary areas that are targeted are the data itself,

whether ciphertext or plaintext, and people, through the use of social engineering techniques. Some of these attack methods are specific to particular types of cryptographic algorithms or systems, while other attack methods are very general in nature. We will discuss many of these attack methods throughout this objective.

Brute Force

A brute-force attack is one in which the attacker has little to no knowledge of the key, the algorithm, or the cryptosystem. Essentially, the attacker is trying every possible combination of ciphertext to derive the correct key (or password), until the correct plaintext is discovered. Most brute-force attacks are offline attacks against password hashes captured from systems through other attack methods, since online attacks are easily thwarted by account lockout mechanisms. Theoretically, given enough computational power, almost all offline brute-force attacks would succeed eventually, but the extraordinary length of time required to break some of the more complex algorithms and keys makes most such attacks infeasible.

It's also important to distinguish a dictionary attack from a brute-force attack. In a dictionary attack, the attacker uses a much smaller range of characters, often compiled into word lists that are hashed and tried against a targeted password hash. If the password hashes match, then the attacker has discovered the password. If they don't match, then the attack progresses to the next word in the word list. Dictionary attacks are often accelerated by using precomputed hashes (called rainbow tables). Note that a dictionary attack is less random than a brute-force attack. In a brute-force attack, the attacker uses all possible combinations of allowable characters to attempt to guess a correct password or key.



EXAM TIP Both dictionary and brute-force attacks are typically automated and can attempt hundreds of thousands of password combinations per minute. Whereas a brute-force attack will theoretically eventually succeed, given enough time and computing power, dictionary attacks are limited to the word lists used and may exhaust all possibilities in the list while never discovering the targeted password.

Ciphertext Only

A ciphertext-only attack is one in which the attacker only has samples of ciphertext to analyze. This type of attack is one of the most common types of attacks since it's very easy to get ciphertext by intercepting network traffic. There are different methods that can be used for a ciphertext-only attack, including frequency analysis, discussed in the upcoming sections.

Known Plaintext

In this type of attack, an attacker has not only ciphertext but also known plaintext that corresponds with it, enabling the attacker to compare the known plaintext with its ciphertext results to determine any relationships between the two. The attacker looks for patterns that

may indicate how the plaintext was converted to ciphertext, including any algorithms used, as well as the key. The purpose of this type of attack is not necessarily to decrypt the ciphertext that the attacker has, but to be able to gather information that the attacker can use to collect additional ciphertext and then have the ability to decrypt it.

Chosen Ciphertext and Chosen Plaintext

In a chosen-ciphertext attack, the attacker generally has access to the cryptographic system and can choose ciphertext to be decrypted into plaintext. This type of attack isn't necessarily used to gain access to a particular piece of encrypted plaintext, but to be able to decrypt future ciphertext. The attacker may also want to discover how the system works, including deriving the keys and algorithms, if the attacker doesn't already have that information.

The chosen-plaintext attack, on the other hand, enables the attacker to choose the plaintext that gets encrypted and view the resulting ciphertext. This also allows the attacker to compare each plaintext input with its corresponding ciphertext output to determine how the cryptosystem works. Again, both of these types of attack assume that the attacker has some sort of access to the cryptographic system in question and can feed in plaintext and obtain ciphertext from the system at will.

Frequency Analysis

Frequency analysis is a technique used when an attacker has access to ciphertext and looks for statistically common patterns, such as individual characters, words, or phrases in that ciphertext. Think of the popular "cryptogram" puzzles you may see in supermarket magazines. This technique typically only works if the ciphertext is not further scrambled or organized into a more puzzling pattern; grouping the ciphertext into distinct groups of ten characters, for example, regardless of how they are spaced apart in the plaintext message, will help defeat this technique.

Implementation

Implementation attacks target not just the key or algorithm but how the cryptosystem in general is constructed and implemented. For example, there may be flaws in how the system stores plaintext in memory before it is encrypted, and this might enable an attacker to access that memory before the encryption process even occurs. Other systems may store decrypted text or even keys in memory. There are a variety of different attacks that can be used against cryptosystems, a few of which we will discuss in the next few sections.

Side Channel

A *side-channel attack* is any type of attack that does not directly attack the key or the algorithm but rather attacks the characteristics of the cryptosystem itself. A side-channel attack may attack different aspects of how the cryptographic system is implemented indirectly. The goal of a side-channel attack is to attempt to gain information on how the cryptosystem works, such as

by recording power fluctuations, CPU processing time, and other characteristics of the cryptosystem, and deriving information about the sensitive inner workings of the cryptographic processes, possibly including narrowing down which algorithms and key sizes are used. There are many different methods of executing a side-channel attack; often the attacker's choice of method depends on the type of cryptosystem involved. Among these methods are fault injection and timing attacks.

Fault Injection

A fault injection attack attempts to disrupt a cryptosystem sufficiently to cause it to repeat sensitive processes, such as authentication, giving the attacker an opportunity to gain information about those processes or, at the extreme, intercept credentials or insert themselves into the process. A classic example of a fault injection attack is when an attacker broadcasts deauthentication traffic over a wireless network and disrupts communications between a wireless client and a wireless access point. This causes both the client and the access point to have to reauthenticate to each other, leaving the door wide open for the attacker to intercept the four-way authentication handshake that the wireless access point uses.

Timing

Timing attacks take advantage of faulty sequences of events during sensitive operations. In Objective 3.4, we discussed how timing attacks exploit the dependencies on sequence and timing in applications that execute multiple tasks and processes concurrently, which includes cryptographic applications and processes. Remember that these attacks are called *time-of-check to time-of-use (TOC/TOU)* attacks. They attempt to interrupt the timing or sequencing of tasks cryptographic systems must execute in sequence. If an attacker can interrupt the sequence or timing, the interruption may allow an attacker to intercept credentials or inject themselves into the cryptographic process. This type of attack is also referred to as an *asynchronous* attack.

Man-in-the-Middle (On-Path)

An on-path attack, formerly known widely as a man-in-the-middle (MITM) attack, can be executed using a couple of different methods. The first method is to intercept the communications channel itself and attempt to use captured keys, certificates, or other authenticator information to insert the attacker into the conversation. Then the attacker can intercept what is being sent and received and has the ability to send false messages or replies. This is related to a similar attack, called a *replay* attack, where the attacker initially intercepts and then "replays" credentials, such as session keys. The second method is to attempt to discover and reverse the encryption and decryption functions to a point that the remaining functions can be determined. This variation is known as a *meet-in-the-middle* attack.

Pass the Hash

In certain circumstances, Windows will use what is known as *pass-through authentication* to authenticate a user to an additional resource. This may happen without the user explicitly having to authenticate, since password hashes are stored in the Windows system even after authentication is complete. In this attack, the perpetrator intercepts the Windows password hash used during pass-through authentication and attempts to reuse the hash on the network to authenticate to other hosts by simply passing the hash to them. Note that the attacker doesn't have the actual password itself, only the password hash. The resource assumes that the password hash comes from a valid user's pass-through authentication attempt.

Pass-the-hash attacks are most often seen in systems that still use older authentication protocols, such as NTLM, but even modern Windows systems can default to using NTLM during certain conditions while communicating with peer systems on a Windows Active Directory network. This makes the pass-the-hash attack still a serious problem even with modern Windows networks.

Kerberos Exploitation

Kerberos, described in detail in Objective 5.6, is highly dependent on consistent time sources throughout the infrastructure, since it timestamps tickets and requests and only allows a very minimal span of time that they can be used. Disrupting or attacking the Kerberos realm's authoritative time source can help an attacker engage in replay attacks. Additional attacks on Kerberos include attempting to intercept tickets for reuse, as well as the aforementioned pass-the-hash attacks. Note that the Kerberos Key Distribution Center (KDC) can also be a single point of failure if compromised.

Ransomware

Ransomware attacks are a twist on traditional cryptographic attacks. Instead of the attacker attempting to steal credentials, discover keys, attack cryptosystems, and decrypt ciphertext, the attacker uses cryptography to hold an organization's data hostage by encrypting it and demanding a ransom in exchange for the decryption key. The attacker frequently threatens the organization by encrypting sensitive and critical data, which the organization cannot get back, and sometimes even threatens to release encrypted data, along with the key, to the public Internet.

Almost all ransomware attacks occur after a malicious entity has invaded an organization's infrastructure through attack vectors, very often phishing attacks. Ransomware has recently been used to attack hospitals, school districts, manufacturing companies, and other critical infrastructure. One of the most recent examples was the Colonial Pipeline attack, showing that ransomware is rapidly becoming the most impactful cyberthreat that organizations face.

REVIEW

Objective 3.7: Understand methods of cryptanalytic attacks In this objective we discussed the various attack methods that can be attempted against cryptographic systems. Cryptanalytic attacks typically target keys, algorithms, and implementation of the cryptosystem itself.

Attacks that target keys include common brute-force and dictionary attacks. Algorithm attacks include more sophisticated attacks that target both pieces of ciphertext and plaintext to gain insight into how the encryption/decryption process works. We also examined frequency analysis and chosen ciphertext attacks. Implementation attacks include side-channel attacks, fault injection attacks, and timing attacks. We also discussed on-path (formally known as man-in-the-middle or MITM) attacks that attempt to interrupt communications or cryptographic processes. Other advanced techniques include pass-the-hash and Kerberos exploitation attacks. We concluded our discussion with how ransomware attacks use cryptography in a different way to deny people the use of their data.

3.7 QUESTIONS

1. Your company has recently endured a cyberattack, and researchers have discovered that many different users' encryption keys were compromised. The post-attack analysis indicates that the attacker was able to hack the application that generates keys, discovering that keys are temporarily stored in memory until the application is rebooted, and was therefore able to steal the keys directly from the application. Which of the following best describes this type of attack?
   - **A.** Fault injection attack
   - **B.** Timing attack
   - **C.** Ransomware attack
   - **D.** Implementation attack
2. You are a security researcher, and you have discovered that a web-based cryptographic application is vulnerable. If it receives carefully crafted input, it can be made to issue valid encryption keys to anyone, including an attacker. Which of the following best describes this type of attack?
   - **A.** Timing attack
   - **B.** Fault injection
   - **C.** Frequency analysis
   - **D.** Pass-the-hash

3.7 ANSWERS

1. **D** Since this type of attack took advantage of a flaw in the application that generates keys, this would be an implementation attack.
2. **B** In this type of attack, a web application could receive faulty input, which creates an error condition and causes valid encryption keys to be issued to an attacker. This would be considered a fault injection attack.



Objective 3.8 Apply security principles to site and facility design

In this objective and the next one, Objective 3.9, we'll discuss physical security elements. First, in this objective, we will explore how the secure design principles outlined in Objective 3.1 apply to the physical and environmental security controls, particularly site and facility design.

Site and Facility Design

Although much of the CISSP exam relates to technical and administrative controls, it also covers physical controls, which are also critically important to protect assets. Site and facility design are advanced concepts that many cybersecurity professionals don't have knowledge of or experience with. We're going to discuss how to secure an organization's site and its facilities and how to design and implement controls to protect people, equipment, and facilities.

Site Planning

If an organization is developing a new site, it has the opportunity to design the premises and facilities to address a wide variety of threats. If the organization is taking over an existing facility, particularly an old one, it may have to make some adjustments in the form of remodeling, retrofitting, landscaping, and so on, to ensure that fundamental security controls are in place to meet threats.

Site planning focuses on several key areas:

- **Crime prevention and disruption** Placement of fences, security guards, and warning signage, as well as implementation of physical security and intrusion detection alarms, motion detectors, and security cameras. Crime disruption (delay) mechanisms are layers of defenses that slow down an adversary, such as an intruder, and include controls of locks, security personnel, and other physical barriers.

- **Reduction of damage** Minimization of damage resulting from incidents, by implementing stronger interior walls and doors, adding external barriers, and hardening structures.
- **Incident prevention, assessment, and response** Assessment and response controls include the same physical intrusion detection mechanisms and security guards. Response procedures include fire suppression mechanisms, emergency response and contingency procedures, law enforcement notification, as well as external medical and security entities.

Common site planning steps to help ensure that physical and environmental security controls are in place include the following:

1. Identifying governance (regulatory and legal) requirements that the organization must meet
2. Defining risk for the physical and environmental security program, which includes assessment, analysis, and risk response
3. Developing controls and countermeasures in response to a risk assessment, as well as their performance metrics
4. Focusing on deterrence, delaying, detection, assessment, and response processes and controls

Cross-Reference

Objective 1.10 covered risk management concepts in depth.

Secure Design Principles

In Objective 3.1, we discussed several key principles that are critical in security architecture and engineering. Security design is important so that systems, processes, and even the security controls involved don't have to be redesigned, reengineered, or replaced at a later date. Although we discussed secure design principles in the context of administrative and technical controls, these same principles apply to physical and environmental design.

Threat Modeling

You already know that threat modeling goes beyond simply listing generic threats and threat actors. For threat modeling to be effective, you must go into a deeper context and relate probable threats with the inherent vulnerabilities in assets unique to your organization. Threat modeling in the physical and environmental context works the same way and requires that you do the following:

- Determine generic threat actors and threats
- Relate those threat actors and threats to the specifics of the organization, its facilities, site design, and weaknesses that are inherent to those characteristics

- Take into account the organization's schedules and traffic flows, as well as its natural and artificial barriers
- Examine the assets that someone would want to physically interact with, to include disrupting them, destroying them, or stealing them

Using threat modeling for your physical environment will help you to design very specific physical and environmental controls to counter those threats.

Least Privilege

As with technical assets and processes, personnel should only have the least amount of physical access privileges assigned to them that are necessary to perform their job functions. Not all personnel should have the same access to every physical area or piece of equipment. While all employees have access to common areas, fewer have access to sensitive processing areas. Least privilege is a physical control when implemented as a facility or sensitive area access control list.

Defense in Depth

The principle of defense in depth (aka layered security) applies to physical and environmental security just as it applies to technical security. Physical security controls are also layered to provide strong protection even when one or more layers fail or are compromised. Layers of physical control include security and safety signage, access control badges, video surveillance and recording systems, physical perimeter barriers, security guards, and the arrangement of centralized entry/exit points.

Secure Defaults

The principle of secure defaults means that security controls are initially locked down and then relaxed only as needed. This includes default physical access to sensitive areas, entrance and exit points, parking lots, emergency exits, nonemergency doors, and storage facility exits and entrances. As an administrative control, the default is to grant access to sensitive areas only to people who need that access to perform their job functions. By default, access is not granted to everyone.

Fail Securely

The term "fail secure" means that in the event of an emergency, security controls default to a secure mode. For example, this can include doors to sensitive areas that lock in the event of a theft or intrusion by unauthorized personnel. Contrast this to the term "fail safe," which means that when a contingency or emergency happens, certain controls fail to an open or safe mode. A classic example is that the doors to a data center should fail to a safe mode and remain unlocked during a fire to allow personnel to escape safely.

Whether to use fail secure controls or fail safe controls is a design choice that management must carefully consider, since the safety and the preservation of human life are the most important aspects of physical security. However, at the same time, assets must be protected from theft, destruction, and damage. A balance may be to implement controls that are complex

and advanced enough to be programmed for certain scenarios and fail appropriately. For example, in the event of an intrusion alarm, either automatic systems or security guards could manually ensure that all doors fail secure, but the same doors, in the event of a fire alarm, may unlock and remain open.



EXAM TIP Note that in the event of an incident, *fail secure* means that controls will fail to a more secure state; however other controls may *fail safe* to a less secure or open state. Fail safe is used to protect lives and ensure safety, while fail secure is used to protect equipment, facilities, systems and information.

Separation of Duties

Just as technical and administrative controls often require a separation of duties, physical controls often necessitate this same principle. Physical security duties are normally separated to mitigate unauthorized physical access, intrusion, destruction of equipment, criminal acts, and other unauthorized activities. For example, separation of duties applied to physical controls could be a policy that requires a person in management to sign in guests but requires another employee to escort the visitors. This demonstrates a two-person control, in that the person who authorizes the visitor is not the same person who escorts them, and adds the assurance that someone else knows they are in the vicinity.

Keep It Simple

The design principle of keep it simple can be applied to physical and environmental security design as well. Simpler physical security design makes the physical layout more conducive to the work environment, reducing traffic and unnecessary movement, and makes maintaining controlled access throughout the facility easier. The simpler facility or workplace design can also help eliminate hiding spots for intruders and help with the positioning of security cameras and guards.

Along with a simpler layout, straightforward procedures to allow access in and throughout the facility are a necessary part of simpler physical design. Overly complicated procedures will almost certainly ensure that those procedures fail at some point, for a variety of reasons:

- Employees don't understand them.
- Employees become too complacent to follow them.
- The procedures interrupt business processes to an unacceptable level.
- The procedures don't meet the functional needs of the organization.

Additionally, simpler procedures for access control within a facility can greatly enhance both security and safety.

Zero Trust

Recall from Objective 3.1 that the zero trust principle means that no entity trusts another entity until that trust has been conclusively authenticated each time there is an interaction between those entities. All entities start out as untrusted until proven otherwise. Physical and procedural controls must be implemented to initially establish trust with entities in the facility and to subsequently verify that trust. These controls include both physical and personnel security measures, such as:

- Positive physical and logical identification of the individual
- Nondisclosure agreements
- Security clearances
- Need-to-know for specific areas
- Supervisor approval for sensitive area access
- Relevant security training

Trust But Verify

As we discussed earlier in the domain, trust is not always automatically maintained after it has been established. Entities can become untrusted, or even compromised. Following the same lines as the zero trust principle, trust must be periodically reverified and the organization must ensure that trust is still valid. There are several ways to validate ongoing trust, and most of these relate to auditing and accountability measures, which may include verifying physical logs, reviewing video surveillance footage, conducting periodic reacknowledgement of the rules of behavior in recurrent security training, and revalidating need-to-know and security clearances from time to time.

Privacy by Design

Privacy by design as a design principle ensures that privacy is considered even before a security control is implemented. In physical environment security planning, facilities and workspaces are designed so that individual privacy is considered as much as possible. Some obvious work areas that must be considered include those where an employee has a reasonable expectation of privacy, like a restroom or locker room. But these also include work areas where sensitive processing of personal data occurs, as well as supervisory work areas where employees are expected to be able to discuss personal or sensitive information with their supervisors. Other spaces that should be considered for privacy include healthcare areas, such as company clinics, and human resources offices.

Shared Responsibility

As discussed in Objective 3.1, the shared responsibility principle of security design means that an organization, such as a service provider, often shares responsibility with its clients. The classic example is the cloud service provider and a client that receives services from that provider.

In the physical and environmental context, however, there is also a paradigm of shared responsibility. Take, for example, a large facility that houses several different companies. Each company may have a small set of offices or work areas that are assigned to them. However, overall responsibility for the facility security may fall to the host organization and could include providing a security guard staff, surveillance cameras, a centralized reception desk with a controlled entry point into the facility, and so on. The tenant organizations in the facility may be responsible for other physical security controls, such as access to their own specific area, and assistance in identifying potential intruders or unauthorized personnel.

REVIEW

Objective 3.8: Apply security principles to site and facility design In this objective we discussed site and facility design, focusing on the secure design principles we discussed first in Objective 3.1. These principles apply to physical and environmental security design in much the same way as they apply to administrative and technical controls. We discussed the need for threat modeling in the physical environment, so physical threats can be understood and mitigated. The principle of least privilege applies to physical space in environments that need to restrict access. Defense-in-depth principles ensure that the physical environment has multiple levels of controls to protect it. We talked about secure defaults for controls that may be configured as more functional than secure, as well as the definitions of fail secure and fail safe. Remember that fail secure means that if a control fails, it will do so in a secure manner. Fail safe applies to those controls that must fail in an open or safe manner to preserve lives and safety. We also discussed separation of duties, and how it applies to the physical environment. The keep it simple principle applied in site and facility design helps to ensure security while not over complicating controls that may interfere with security or safety. Zero-trust people are not trusted by default with access to physical facilities; they must establish that trust and maintain it throughout their time in the facility. Additionally, we discussed the trust but verify principle, meaning that trust must be occasionally reestablished and verified as well. Privacy by design ensures that private spaces are included in site and facility planning. Finally, shared responsibility addresses facilities where there may be multiple organizations that share the facility and security functions must be shared amongst them.

3.8 QUESTIONS

1. Which of the following principles states that personnel should have access only to the physical areas they need to enter to do their job, and no more than that?
   - **A.** Separation of duties
   - **B.** Least privilege
   - **C.** Zero trust
   - **D.** Trust but verify

2. In your company, when personnel first enter the facility, they not only must swipe their electronic access badge in a reader, which verifies who they are, but also must pass through a security checkpoint where a guard visually verifies their identification by viewing the picture on their badges. Periodically throughout the day, they must swipe their access badges in additional readers and may be subject to additional physical verification. Which of the following principles is at work here?
   - **A.** Trust but verify
   - **B.** Separation of duties
   - **C.** Privacy by design
   - **D.** Secure defaults

3.8 ANSWERS

1. **B** The principle of least privilege provides that personnel in an organization should have access only to the physical areas that they need to enter to perform their job functions, and no more than that. This ensures that people do not have more access to the facility and secure work areas than they need.
2. **A** Since personnel must initially verify their identity when entering the facility and then periodically reverify it throughout the day, these actions conform to the principle of trust but verify.



Objective 3.9 Design site and facility security controls

In this objective we're continuing our discussion of physical and environmental security control design. We're moving beyond the basic principles of site and facilities security that were introduced in Objective 3.8 to cover key areas and factors you must address during the building and floorplan design processes.

Designing Facility Security Controls

Designing security controls for the physical environment requires that you understand not only the design principles we discussed in Objectives 3.1 and 3.8, but also how they may apply to different situations, site layouts, and facility characteristics. You also need to understand the security requirements of your organization. As mentioned in the previous objectives, site and facility design require that you consider key goals, such as preventing crime, detecting events, protecting personnel and assets, and delaying the progress of intruders or other malicious actors until security controls can be invoked, all while assessing the situation. Achieving these

key goals requires you to consider some key areas of focus you should take into consideration, which we will discuss in the upcoming sections, as well as specifically how to prevent criminal or malicious acts through purposeful environmental design.

Crime Prevention Through Environmental Design

Preventing crime or other malicious acts is the focus of a discipline called Crime Prevention Through Environmental Design (CPTED). The theory behind CPTED is that the environment that human beings interact with can be manipulated, designed, and architected to discourage and reduce malicious acts. CPTED addresses characteristics of a site or facility such as barrier placement, lighting, personnel and vehicle traffic patterns, landscaping, and so on. The idea is to influence how people interact with their immediate environment to deter malicious acts or make them more difficult. For example, increasing the coverage of lighting around a facility may discourage an intruder from breaking into the building at night; the placement of physical barriers, such as bollards, helps to prevent vehicles from charging an entryway of a building.

CPTED addresses the following four primary areas:

- **Natural access control** This entails naturally guiding the entrance and exit processes of a site or facility by controlling spaces and the placement of doors, fences, lighting, landscaping, sidewalks, and other barriers.
- **Natural surveillance** Natural surveillance is intended to make malicious actors feel uncomfortable or deterred by designing environmental features such as sidewalks and common areas to be highly visible so that observers (other than guards or electronic means) can watch or surveil them.
- **Territorial reinforcement** This is intentional physical site design emphasizing or extending an organization's physical sphere of influence. Examples include walls or fencing, signage, driveways, or other barriers that might show the property ownership or boundary limits of an organization.
- **Maintenance** This refers to keeping up with physical maintenance to ensure that the site or facility presents a clean, uncluttered, and functional appearance. Repairing broken windows or fences, as well as maintaining paint and exterior details, demonstrate that the facility is cared for and well kept, and shows potential intruders that security is taken seriously.

Key Facility Areas of Concern

In every site or facility, several areas warrant special attention when designing and implementing security controls. These include restricted work areas, server rooms, and large data centers. Storage facilities, such as those used to store sensitive media or investigative evidence, also must be carefully controlled. Designing site and facility security controls encompasses not

only protecting data and assets, but also ensuring personnel safety. Safety features that must be carefully planned and implemented include appropriate utilities and environmental controls as well as fire safety and electrical power management.

Wiring Closets/Intermediate Distribution Facilities

Any organization that relies on an Internet service provider (ISP) to deliver high-bandwidth communications services has specific sensitive areas and equipment on its premises that are devoted to receiving those services. These communications areas are called *distribution facilities* and are the physical points where external data lines come into the building and are broken out from higher bandwidth lines into several lower bandwidth lines. In larger facilities, these are called *main distribution facilities*, or MDFs. Smaller distribution areas and equipment are referred to as *intermediate distribution facilities*, or IDFs. IDFs break out high-bandwidth connections into individual lines or network cabling drops for several endpoints, either directly to hosts or to centralized network switches. MDFs are usually present in data centers and server rooms of large facilities, whereas IDFs commonly are located in small wiring closets.

Usually, MDFs are secured because they are located in large data centers or server rooms that are almost always protected. IDFs may sometimes be less secure and located in unprotected common areas or even in janitors' closets. Ideally, both MDFs and IDFs should be in locked or restricted areas with limited access. They should also be elevated in the event of floods or other types of damage, so they obviously should not be implemented in basement areas or below ground level. They should also be located away from risks presented by faulty overhead sprinklers, broken water pipes, and even heating, ventilation, or air conditioning (HVAC) equipment.

Server Rooms/Data Centers

Since server rooms and large data centers contain sensitive data and processing equipment, these areas must be restricted only to personnel that have a need to be there to perform their job function. In terms of facility design, it's better to locate data centers, server rooms, and other sensitive areas in the core or center of a facility, to leverage the protection afforded by areas of differing sensitivity, such as thick walls. Additionally, server rooms and data centers should be located in the lower floors of a building so that emergency personnel will have easier access.

There are specific safety regulations that must be implemented for critical server rooms or data centers. Most fire regulations require at least two entry and exit doors in controlled areas for emergencies. Server room and data center doors should also be equipped with alarms and fail to safe mode (fail safe) in the event of an emergency.

Cross-Reference

The use of "fail safe" mechanisms to preserve human health and safety was discussed in Objective 3.1.

Additional security characteristics of server rooms and data centers include

- Positive air pressure to ensure smoke and other contaminants are not allowed into sensitive processing areas
- Fire detection and suppression mechanisms
- Water sensors placed strategically below raised floors in order to detect flooding
- Separate power circuits from other equipment or areas in the facilities, as well as backup power supplies

Media Storage Facilities

Media should be stored with both security and environmental considerations in mind. Due to the sensitive nature of data stored on media, the media should be stored in areas with restricted access and environmental controls. Media with differing sensitivity or classification levels should be separated. Note that depending upon what type of backup media are present, such as tape or optical media, the facility operators must carefully control temperature and humidity in media storage locations. All types of media, regardless of sensitivity or backup type, should be under inventory control so that missing sensitive media can easily be identified and located.

Evidence Storage

Evidence storage has its own security considerations, but sensitive areas within facilities that are designated as evidence storage should also have at least the following considerations:

- Highly restrictive access control
- Secured with locks and only one entrance/exit
- Strict chain of custody required to bring evidence in or out of storage area
- Segregated areas within evidence storage that apply to evidence of differing sensitivity levels

Restricted and Work Area Security

All sensitive areas, whether designated for processing sensitive data or storage, must be designed with security and safety in mind. Sensitive work areas should have restrictive access controls and only allow personnel in the area who have a valid need to be there to perform their job functions. Sensitive work areas need to be closely monitored using intrusion detection systems, video surveillance, and so on. Strong barriers, as well as locking mechanisms and wall construction, should be implemented whenever possible. Work areas containing data of different sensitivity levels should be designed so that they are segmented by need-to-know. Both dropped ceilings and raised floors, when implemented, should be equipped with alarms and carefully monitored to prevent intruders.

Utilities and Heating, Ventilation, and Air Conditioning

Facility utilities include electric power, water, communications, and heating, ventilation, and air conditioning (HVAC) services. HVAC is required to ensure proper temperature and humidity control in the processing environment, and requires redundant power to ensure this environment is properly maintained. Sensitive areas need positive drains (meaning content flows out) built into floors to protect against flooding. Additionally, water, steam, and gas lines should have built-in safety mechanisms such as monitors and shut-off valves. Spaces should be designed so that equipment can be positioned within processing areas with no risk that a burst or leaking pipe will adversely affect processing, damage equipment, or create unsafe conditions.

Environmental Issues

Depending upon the geographical location of the facility, wet/dry or hot/cold climates can create issues by having more or less moisture in the air. Areas with high humidity and hotter seasons may have more moisture in the air that can cause rust issues and short-circuits. Static electricity is an issue in dry or colder climates due to less moisture in the air. Static electricity can cause shorts, seriously damage equipment, and create unsafe conditions. In addition to monitoring humidity and temperature controls, use of thermometers/thermostats to monitor and modify temperature and hygrometers to control humidity necessary. A hygrothermograph can be used to measure and control both temperature and humidity concurrently.

Fire Prevention, Detection, and Suppression

Addressing fires involves three critical efforts: preventing them, detecting them, and suppressing or extinguishing them. Of these three areas, fire prevention is the most important. Fire prevention is accomplished in several ways, including training personnel on how to prevent fires and, most importantly, how to react when fires do occur.

Another critical part of fire prevention is having a clean, uncluttered, and organized work environment to reduce the risk of fire. Construction and design of the facility can contribute to preventing fires; walls, doors, floors, equipment cabinets, and so on should have the proper fire resistance rating, which is typically required by building or facility codes. Fire resistance ratings for building materials are published by ASTM International.

Fires can be detected through both manual and automatic detection means. Obviously, people can sense fires through sight or smell and report them, but often this might be slower than other detection mechanisms. Smoke-activated sensors can sound an alarm before the fire suppression system activates. These sensors are typically photoelectric devices that detect variations in light intensity caused by smoke. There are also heat-activated sensors that can sound a fire alarm when a predefined temperature threshold is reached or when temperature increases rapidly over a specific time period (called the *rate of rise*). Fire detection sensors should be placed uniformly throughout the facility in key areas and tested often.

TABLE 3.9-1 Combustion Elements and Fire Suppression Methods

Combustion Element	Suppression Element	How Suppression Works
Fuel	Soda acid	Removes fuel
Oxygen	Carbon dioxide	Displaces oxygen
Temperature	Water	Reduces temperature
Chemical combustion	Gas (non-halon)	Halts the chemical reaction between elements



NOTE Rate-of-rise sensors provide faster warnings than fixed-temperature sensors but can be prone to false alarms.

Fire requires three things: a fuel source, oxygen, and an ignition source. Reducing or taking away any one of these three elements can prevent and stop fires. Suppressing a fire means removing its fuel source, denying it of oxygen, or reducing its temperature. Table 3.9-1 lists the elements of fires and how they may be extinguished or suppressed.

Fire suppression also includes having the right equipment on hand close by to extinguish fires by targeting each one of these elements. Fire suppression methods should be matched to the type of fire, as well as its fuel and other characteristics. Table 3.9-2 identifies the U.S. classification of fires, including their class, type, fuel, and the suppression agents used to control them.

TABLE 3.9-2 Characteristics of Common Fires and Their Suppression Methods

Class	Type	Fuel	Suppression/ Extinguishing Agent
A	Common combustibles	Wood products, paper, laminates	Water, foam, dry powders, wet chemicals
B	Liquids and gases	Petroleum products and coolants, butane, propane, methane	CO_2, foam, dry powders
C	Electrical	Electrical equipment and wiring	CO_2, dry powders
D	Metals	Aluminum, lithium, magnesium	Dry powders
K	Cooking oils and fats	Kitchens/break rooms	Wet chemicals



CAUTION Using the incorrect fire suppression method not only will be ineffective in suppressing the fire, but may also cause the opposite effect and spread the fire or cause other serious safety concerns. An example would be throwing water on an electrical fire, which could create a serious electric shock hazard.



EXAM TIP You should be familiar with the types and characteristics of common fire extinguishers for the exam.

There are some key considerations in fire suppression that you should be aware of; most of these relate to the use of water suppression systems, since water is not the best option for use around electrical equipment, particularly in data centers. However, water pipe systems are still used throughout other areas of facilities. Water pipe or sprinkler systems are usually much simpler and cost less to implement. However, they can cause severe water damage, such as flooding, and contribute to electric shock hazard.

There are four main types of water sprinkler systems:

- **Wet pipe** This is the basic type of system; it always contains water in the pipes and is released by temperature control sensors. A disadvantage is that it may freeze in colder climates.
- **Dry pipe** In this type of system there is no water kept in the system; the water is held in a tank and released only when fire is detected.
- **Preaction** This is similar to a dry pipe system, but the water is not immediately released. There is a delay that allows personnel to evacuate or the fire to be extinguished using other means.
- **Deluge** As its name would suggest, this allows for a large volume of water to be released in a very short period, once a fire is detected.



EXAM TIP You should be familiar with the four main types of water sprinkler systems: wet pipe, dry pipe, preaction, and deluge.

There are some other key items you should remember for the CISSP exam, and in the real world, when dealing with fire prevention, detection, and suppression. First, you should ensure that personnel are trained on detecting and suppressing fires and, more importantly, that there is an emergency evacuation system in place in the event personnel cannot control the fire.

This evacuation plan should be practiced frequently. Second, HVAC systems should be connected to the alarm and suppression system so that they shut down if a fire is detected. The HVAC system could actually spread the fire by supplying air to it as well as conveying smoke throughout the facility. Third, since cabling is often run in the spaces above dropped ceilings, you should ensure that only plenum-rated cabling is used. This means that the cabling should not be made of polyvinyl chloride (PVC), since burning those types of cables can release toxic gases harmful to humans.

Power

All work areas in the facility, especially data centers and server rooms, require a constant supply of clean electricity. There are several considerations when dealing with electrical power, including backup power and power fluctuations. Note that redundant power is applied to systems at the same time as main power; there is no delay if the main power fails, and redundant power takes over. Backup power, on the other hand only comes on after a main power failure. Backup power strategies include

- **Uninterruptible power supplies** Used to supply battery backup power to an electrical device, such as a server, for only a brief period of time so that the server may be gracefully powered down.
- **Generators** Self-contained engine that runs to provide power to a facility for a short-term (typically hours or days), which must be fueled on a frequent basis. Note that generators can quickly produce large volumes of dangerous exhaust gases and should only be operated by trained personnel.

Power issues within a facility can be short or long term. Even momentary power issues can cause damage to equipment, so it's vitally important that power be carefully controlled and conditioned as it is supplied to the facility and the equipment. Power issues can include a momentary interruption or even a momentary increase in power, either of which can damage equipment.

Power loss conditions include *faults*, which are momentary power outages, and *blackouts*, which are prolonged and usually complete losses of power. Power can also be degraded without being completely lost. A *sag* or dip is a momentary low-voltage condition usually only lasting a few seconds. A *brownout* is a prolonged dip that is below the normal voltage required to run equipment. Momentary increases of power include an *inrush* current condition, which is an initial surge of current required to start a load, which usually happens during the switchover to a generator, and also a *surge* or *spike*, which is a momentary increase in power that may burn out or otherwise damage equipment. Voltage regulators and line conditioners are electrical devices connected inline between a power supply and equipment to ensure clean and smooth distribution of power throughout the facility, data center, or perhaps just a rack of equipment.

REVIEW

Objective 3.9: Design site and facility security controls In this objective we completed our discussion of designing site and facility security controls using the security principles we covered in Objectives 3.1 and 3.8. This discussion applied those principles to security control design and focused on crime prevention through the purposeful design of environmental factors, such as lighting, barrier placement, natural access control, surveillance, territorial reinforcement, and maintenance. We also discussed protection of key facility areas, such as the main and intermediate distribution facilities, which provide the connection to external communications providers and distribute communication service throughout the facility. We considered the safety and security of server rooms and data centers, as well as media storage facilities, evidence storage, and sensitive work areas. We talked about security controls related to utilities such as electricity, communications, and HVAC. We touched upon the need to monitor critical environmental issues such as humidity and temperature to ensure that they are within the ranges necessary to avoid equipment damage. We also covered the importance of fire prevention, detection, and suppression, and how those three critical processes work. Finally, we assessed power conditions that may affect equipment and some solutions to minimize impact.

3.9 QUESTIONS

1. You are working with the facility security officer to help design physical access to a new data center for your company. Using the principles of CPTED, you wish to ensure that anyone coming within a specific distance of the entrance to the facility will be easily observable by employees. Which of the following CPTED principles are you using?
   - **A.** Natural surveillance
   - **B.** Natural access control
   - **C.** Maintenance
   - **D.** Territorial reinforcement
2. Which of the following types of fire suppression methods would be appropriate for an electrical fire that may break out in a server room?
   - **A.** This is a Class A fire, so water or foam would be appropriate.
   - **B.** This is a Class B fire, so wet chemicals would be appropriate.
   - **C.** This is a Class K fire, so wet chemicals would be appropriate.
   - **D.** This is a Class C fire, so CO_2 would be appropriate.

3.9 ANSWERS

1. **A** In addition to electronic surveillance measures, you want to design the physical environment to facilitate observation of potential intruders or other malicious actors by normal personnel, such as employees. This is referred to as natural surveillance.
2. **D** An electrical fire is a Class C fire, which is normally suppressed by using fire extinguishers using CO_2 or dry powders. Water, foam, or other wet chemicals would be inappropriate and may create an electrical shock hazard.



Communication and Network Security

Domain Objectives

- **4.1** Assess and implement secure design principles in network architectures.
- **4.2** Secure network components.
- **4.3** Implement secure communication channels according to design.

Domain 4 covers secure networking infrastructures. Secure networking has always been a critical part of the CISSP exam, but over the past few versions of the exam, this domain has shifted from merely requiring memorization of foundational network knowledge, such as network architectures, port numbers, security devices, secure protocols, and so on, to an emphasis on applying all of this knowledge to secure an infrastructure using the secure design principles we discussed in Domain 3. Foundational network knowledge is still critical, and we will cover the key points you need to know for the exam objectives; however, you should focus on how each of these simple network components is used to create a layered approach to strong network security. In this domain we will address three objectives that focus on implementing secure design principles in network architectures, configuring secure network components, and securing network communications channels.



Objective 4.1 Assess and implement secure design principles in network architectures

This objective covers fundamental security design principles as they are applied to network architectures. To understand the application of these principles, we will review the key fundamentals of networking. While we will not cover networking to a great depth, we will review core concepts such as the OSI model and TCP/IP stack, IP networking, and secure protocols. We will also discuss the application of networking technologies that have evolved over the years, such as multilayer protocols, converged technologies, microsegmentation, and content distribution networks. We will not be limiting the discussion to wired networks, as we will also review the key points of wireless technologies, including Wi-Fi and cellular networks.

Fundamental Networking Concepts

Understanding the fundamental networking concepts covered in this section not only is necessary for passing the CISSP exam but also will help you to understand how to apply more complex security design principles to network architectures. In this section we're going to cover both the Open Systems Interconnection (OSI) and Transmission Control Protocol/Internet Protocol (TCP/IP) models, which help to conceptualize how networks function at different layers, using various protocols and devices, as well as how data is passed between layers. Network communications models such as these help network professionals such technicians, engineers, and architects conceptualize the flow of data across the network, as well as design and maintain infrastructures in a secure manner. In addition to discussing the OSI model and TCP/IP stack, we will also look at secure protocols, since their use is of critical importance in transmitting and receiving data while ensuring confidentiality and integrity.

Open Systems Interconnection and Transmission Control Protocol/Internet Protocol Models

Network communications models are used to demonstrate how network components operate and interact with each other. These include both physical components, such as network devices, and logical components, such as data. A model can help you understand the various functions of different types of equipment, protocols, and data that all interact to facilitate network communications. Keep in mind that models are just abstract representations and are not necessarily representative of all networks. They are designed to help you understand how networks communicate and to provide standards to use when building network architectures, especially with different components.

OSI Model

The OSI model is the ubiquitous standard with which networks are designed and function. In modern networking, components such as protocols, interfaces, and devices are all designed to be interoperable with the OSI model. The OSI model is not a protocol stack or component itself but provides a framework to use for building networks and connecting their components.

The OSI model consists of seven layers, each numbered from 1 to 7, starting with layer 1 at the bottom layer and continuing to layer 7 at the top. Each layer of the model represents the different interactions that happen with network traffic, protocols, devices, and so on at that layer. Each layer represents a different function of networking. Figure 4.1-1 shows the seven layers of the OSI model.

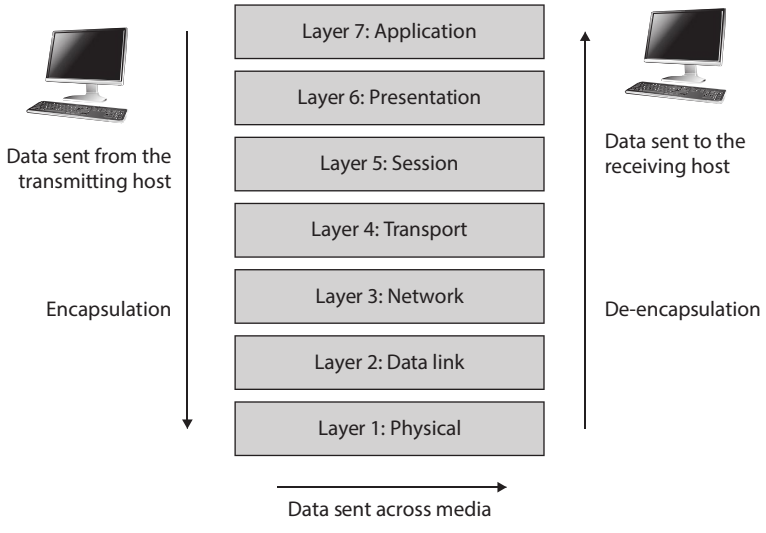


FIGURE 4.1-1 The OSI model

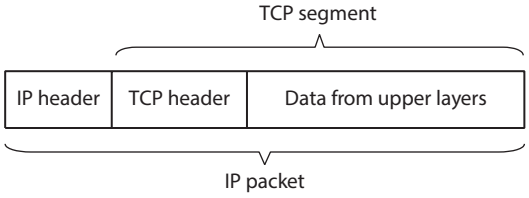


FIGURE 4.1-2 Example of PDUs and headers as data travels between the transport and network layers of the OSI model

Each layer of the model interacts with data differently; however, each layer receives data and passes it to the next layer, either "up" or "down" the model. Each layer ignores all layers except the layers immediately above and below it. Data is transformed within each layer, and is referred to as different protocol data units (PDUs) during this transformation, depending on the layer in which it resides.

As shown in Figure 4.1-1, data passing *down* layers (from 7 to 1) is *encapsulated*, which means that each layer adds administrative header information to the data, which becomes part of the data itself as it is passed down from layer to layer. Conversely, data passing "up" the model, from layer 1 to 7, is *de-encapsulated*, meaning that header information is stripped from each PDU and the remaining data is passed up to the next layer. Figure 4.1-2 shows how header information is added to data coming down from layer 4 (TCP) encapsulated in the next layer, layer 3 (IP).

Different devices and protocols work at various layers of the OSI model; Table 4.1-1 summarizes the OSI model layers, their relevant devices and protocols, and the protocol data units you should remember for the exam.



EXAM TIP Many of these protocols and devices also function at other layers, performing a particular function or interacting with other components in specific ways, so you may see that SSH, for example, also works at the session layer. For the purpose of the CISSP exam, however, you should focus on the primary layer at which the protocol device functions. Protocols which span multiple layers are called *multilayer protocols* and are discussed later in this objective.

TCP/IP Stack

TCP/IP is a suite of protocols (sometimes called a *protocol stack*) consisting of protocols that are implemented based on the OSI model. TCP/IP is often referred to as a model, and that's not necessarily incorrect, but compared to the OSI model, which only exists as a theoretical framework, the TCP/IP stack is an actual set of protocols that function and work to facilitate

TABLE 4.1-1 Protocols, Devices, and PDUs at Various Layers of the OSI Model

Layer	Protocol, Service, or Standard	Devices	PDU
7 – Application	HTTP, DNS, DHCP, SSH	Gateways, servers, workstations, firewalls	Data
6 – Presentation	ASCII, EBCDIC	—	Data
5 – Session	SQL, SIP, PTP	—	Data
4 – Transport	TCP, UDP	Gateways, firewalls	TCP segments or UDP datagrams
3 – Internet	IP, ICMP, IGMP	Routers, layer 3 switches	IP packets
2 – Data link	IEEE 802.3, IEEE 802.11	Switches, bridges	Frames
1 – Physical	Signaling standards	Hubs, cables	Electrical pulses of binary digits (bits and bytes)

network communications every day. In other words, TCP/IP is a practical implementation of the OSI model, not simply a model itself.

TCP/IP consists of four layers. These layers don't correspond on a one-to-one basis to the OSI model, but all of their functionality can be matched directly to OSI. We talk about common protocols such as HTTP, FTP, DNS, IP Security (IPSec), and so on; these protocols are part of the TCP/IP suite. IT professionals often refer to the TCP/IP stack and OSI model interchangeably, since the protocols of the TCP/IP suite, by extension, work at the various layers of the OSI model. In common discussion this is not incorrect, but for the purposes of the exam you need to know the TCP/IP layers and how they correspond to the layers of the OSI model, as shown in Figure 4.1-3.

Internet Protocol Networking

The Internet Protocol, or IP, is one of the critical protocols in the TCP/IP suite. IP provides routing services for network traffic; without it, the traffic would be limited to local area networks only. IP is responsible for logical addressing and routing that carries network data over larger networks, such as the Internet, across multiple WAN links and routers. IP is composed of several different protocols, a few of which we will discuss briefly here. One of the more important protocols in IP, IPSec, is discussed in the next section, however.

IP version 4 (IPv4) is the current version of IP used worldwide. It's the version that has become the backbone protocol of the modern Internet. IPv4 uses 32-bit addresses, which, theoretically, were exhausted a few years ago. However, with the advent of network address translation (NAT), private IP addressing, and other technologies, IPv4 has had its life extended.

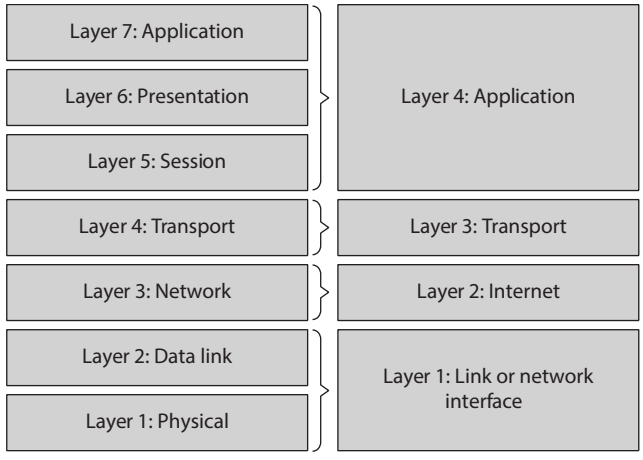


FIGURE 4.1-3 TCP/IP suite mapping to the OSI model

IP version 6 (IPv6) is the next generation of IP and has been adopted on a limited basis throughout the worldwide Internet. Version 6 expands the limited 32-bit addresses of IPv4 to 128 bits used for addressing. It also includes features such as address scoping, autoconfiguration, security, and quality of service (QoS).

IPv4 includes several protocols, such as ICMP, IGMP, and ARP, discussed next. IPSec, a later addition to IPv4, is discussed in the next section in the context of secure protocols.

ICMP

The Internet Control Message Protocol (ICMP) is used to communicate with hosts to determine their status. You can determine if a host is online and whether or not its TCP/IP stack is functioning at some level by using utilities that use ICMP. The most common utilities that network professionals are familiar with are *ping*, *traceroute*, and *pathping*. These utilities use echo requests to send short maintenance messages to a host or network and receive echo replies to respond to the requests, indicating whether or not the host is online. Unfortunately, ICMP can also be used to conduct various network-based attacks on a network or host, although most of these attacks have been mitigated by modern updates to operating systems and the TCP/IP protocol stacks that are installed on them. ICMP inbound to a network is often blocked or filtered for this reason.

IGMP

Internet Group Management Protocol (IGMP) is used to support *multicasting*, which is the process of transmitting data to a specific group of hosts. IGMP uses a specific IP address class, the Class D range, which begins with the 224.0.0.0 address space. Hosts in a particular IGMP group have a normal IP address that is reachable by any other host on the network, as well as an IGMP address in that range.

ARP

Address Resolution Protocol (ARP) is used to resolve a logical 32-bit IP address to a 48-bit hardware address (the physical or MAC address of the network interface for the host). This is because at layer 2, the data link layer of the OSI model, the host only understands hardware addresses for the local network. Before sending data out from the host, the destination logical IP address must be converted to a hardware address and then sent over the transmission media. ARP fulfills this function. If ARP does not resolve an IP address to a local hardware address, the host assumes that the address is on a remote system and forwards it to the hardware address of the default gateway, normally the router interface for the local network.

ARP has security issues that stem from *ARP poisoning*, which means that the local host's ARP cache can be polluted with incorrect entries. This can happen due to gratuitous (unsolicited) ARP requests and false replies that a malicious entity may send out. This can cause an unsuspecting victim host to communicate with a malicious host instead of the one intended.

Secure Protocols

In the early days of the Internet, most protocols were not used to protect data in transit; protocols such as Telnet or HTTP do not natively provide authentication or encryption services, which allows data to be sent in clear text and easily intercepted. Modern protocols, however, ensure that not only is data encrypted, strong authentication and integrity mechanisms are also built in. In this portion of the objective we will discuss several of these secure protocols.

Secure protocols can be used to directly protect data and provide authentication and integrity services, or they can encapsulate unprotected data and nonsecure protocols, giving protection by providing tunneling services for them.

SSL, TLS, and HTTPS

Hypertext Transfer Protocol (HTTP) is a ubiquitous protocol used to transfer web content, such as text, images, and other multimedia. It is the language that creates web pages. However, by itself it is not secure and sends information in clear text. It must use secure network transmission protocols for authentication and encryption. One of these protocols in widespread use is Transport Layer Security (TLS). TLS was always intended as a replacement for an older protocol, Secure Sockets Layer (SSL) but, in 2014, SSL 3.0 was found to be vulnerable to the POODLE attack, which affects all block ciphers that SSL uses. Both protocols operate in similar manners, however, and when HTTP is tunneled through SSL or TLS, it is known as HTTPS (for HTTP Secure). Both SSL and TLS use TCP port 443 by default.



CAUTION Secure Sockets Layer in all its versions (through 3.0) has been deprecated since 2015 and effectively replaced by Transport Layer Security. However, you may still see references to it on the exam, as well as the ability to configure SSL in legacy applications.

TLS, currently in version 1.3 (as of August 2018), provides end-to-end encryption services and works primarily at the session layer of the OSI model. Version 1.3 supports only five encryption algorithms (as opposed to 37, including some that had known vulnerabilities, in previous versions), so this minimizes the number of vulnerable cipher suites that can be used. This also makes it difficult for an attacker to downgrade the level of encryption to a less secure version. TLS not only is used to secure normal HTTP sessions but can also be used for other secure end-to-end encryption needs, such as virtual private networking. TLS-based virtual private networks (VPNs) carry less overhead and can be used for client-to-site VPNs over modern web browsers. TLS supports both server authentication and mutual authentication between the client and server.

When a client initiates a TLS connection, it sends a "hello" message that lists the cipher suites the client supports and a request for key exchange. The server replies with its choice of cipher suites and secure protocols. The server also sends its digital certificate, which proves the server's identity, and provides a public key for the client to use for secure key exchange. The client then sends back a secure session key by encrypting it with the server's public key, which only the server can decrypt. This establishes a secure key for encrypted communications. An optional step is having the client authenticate itself to the server by sending its own digital certificate and public key.

IPSec

IP Security (IPSec) is a protocol suite that resides at the network layer of the OSI model (or the Internet layer of TCP/IP). It was developed because IPv4 does not have any built-in security mechanisms. IPSec can provide both encryption and authentication services, as well as secure key exchange, protecting IP traffic. IPSec consists of four main protocols in the suite:

- **Authentication Header (AH)** Provides for authentication services and data integrity
- **Encapsulating Security Payload (ESP)** Provides encryption services for data payloads
- **Internet Security Association and Key Management Protocol (ISAKMP)** Allows security key association and secure key exchange
- **Internet Key Exchange (IKE)** Assists in key exchange between two entities using IPSec



NOTE IPSec can use AH and ESP separately or together, depending on whether the traffic needs to be encrypted, authenticated, or both.

IPSec can be used to protect communications in one of two modes: transport mode or tunnel mode. In transport mode, IPSec can be used on the local network and can encrypt specific traffic, including specific protocols between multiple hosts, as long as all hosts support the same authentication and encryption methods. This can help secure sensitive traffic

within a network. In tunnel mode, IPSec is tunneled into another protocol, most commonly into the Layer 2 Tunneling Protocol (L2TP), and sent outside of local networks across other, nonsecure or public networks, including the Internet, to another network, making it effective for establishing VPN connections. IPSec can protect both data and header information while in tunnel mode.



EXAM TIP Understand the difference between IPSec's transport and tunnel modes. Transport mode only protects the IP payload and is normally used on an internal network, while tunnel mode encapsulates the entire packet, including header information (e.g., IP address), and is used to carry traffic securely across untrusted networks, such as the Internet. IPSec traffic in tunnel mode must also be encapsulated in a network tunneling protocol such as L2TP.

Secure Shell

Secure Shell (SSH) is both a secure protocol and a suite of tools used to provide encryption and authentication services to communications sessions. It's most commonly used in Linux environments, although it has been ported to Windows operating systems as well. It is often used for remote access between hosts to perform secure, privileged operations. This makes SSH ideal to replace previously used remote session protocols, such as Telnet, which offered no security services whatsoever. Although not very scalable, SSH can be used for long-haul remote access on a limited basis. In addition to encrypting data and providing authentication between hosts, SSH can be used to protect nonsecure protocols, such as File Transfer Protocol (FTP), when those protocols are tunneled through it. However, SSH offers its own secure versions of these protocols, such as Secure Copy Protocol (SCP) and Secure Shell FTP (SFTP).



EXAM TIP Don't confuse SFTP, the Secure Shell version of FTP, with FTP that is simply *tunneled* through TLS or SSH (known as FTPS). SFTP is not tunneled through SSH—it is actually part of the SSH secure suite.

SSH uses TCP port 22 and can work at several layers of the OSI model, including the session layer and higher. Utilities included in the SSH suite allow users to generate host and user keys to facilitate strong authentication between hosts and users.

EAP

The Extensible Authentication Protocol (EAP) is a framework that allows multiple types of authentication methods to be used for users and devices to authenticate networks, typically over remote or wireless connections. EAP can use many different authentication methods, such as

passwords, tokens, biometrics, one-time passwords (OTPs), Kerberos, digital certificates, and several others. When two entities connect using EAP, they negotiate a list of authentication methods that are common to both entities. EAP can be used over a variety of other protocols, such as PPP, PPTP, and L2TP, and over both wired and wireless networks. There are several different variants of EAP, which use TLS (EAP-TLS), pre-shared keys (EAP-PSK), tunneled TLS (EAP-TTLS), and version 2 of the Internet Key Exchange (EAP-IKE2). These variants are each suitable for different types of authentication requirements, depending on the existing infrastructure and compatibility with legacy devices. EAP is also used extensively with 802.1X authentication methods, discussed next.

802.1X

IEEE 802.1X is not part of the 802.11 standards but is often confused with them because it is frequently used in conjunction with WLANs. 802.1X is a port-based authentication protocol and can be used with both wired and wireless networks. Its primary purpose is access control. It allows for both devices and users to be authenticated and can enforce mutual authentication. It is more often encountered in enterprise-level networks than in personal networks.

802.1X has three important components you should be aware of for the exam:

- **Supplicant** Typically a client device
- **Authenticator** Infrastructure device connecting the client to the network and passing on authentication requests to the authentication server
- **Authentication server** Server containing authentication information, such as a RADIUS or Kerberos server

802.1X can use a variety of authentication methods, including the Extensible Authentication Protocol and its many variants (PEAP, EAP-TLS, and EAP-TTLS, among others).

Kerberos

Kerberos is a secure protocol used to authenticate users to networks and resources. Kerberos is most commonly used in Lightweight Directory Access Protocol (LDAP) networks as a single sign-on (SSO) technology. Kerberos is the authentication protocol of choice in Windows Active Directory networks and uses a ticket-based system to authenticate users and then provide authentication services between users and resources. It is also heavily time-based, to prevent replay attacks. Kerberos uses TCP port 88.

Cross-References

Some of these secure protocols, such as SSH, 802.1X, EAP, and TLS, are discussed further in Objective 4.3.
Kerberos will be discussed in greater detail in Objective 5.6, in the context of implementing authentication systems.

Application of Secure Networking Concepts

Elements such as port, protocol, and service are central in secure networking. These elements, along with IP addresses, hardware addresses, and other characteristics of network traffic, are important in securing networks because they are used to *filter* (block or deny) traffic.

Filtering traffic based on port, protocol, and so on only secures the network against some of the heavier "noise" that is relatively easy to block. To perform traffic filtering of a higher fidelity, network devices such as firewalls must be able to examine not only the characteristics of the traffic but the traffic itself. In other words, the content of the traffic must be examined beyond simply looking at the protocol it's using or IP address it has.

Cross-Reference

We will discuss firewalls and other specialized security devices in Objective 7.7.

Beyond traffic filtering based on characteristic or content, secure networking also involves the design and implementation of secure architectures. This means that physical and logical network topologies must be designed with security in mind and use networking components such as security devices, secure protocols, controlled traffic flows, strong encryption, and authentication to create a multilayered approach to protecting systems and information. Secure network architectures are discussed in greater detail in the next objective, 4.2.

Implications of Multilayer Protocols

The term *multilayer protocols* typically refers to either of the following, depending on the context:

- Protocols that span multiple layers of the OSI model or TCP/IP stack. SSH is one example, since it operates at the application and session layers of the OSI model.
- Protocols that operate at the same layer of the OSI model or TCP/IP stack but use both UDP and TCP. The Domain Name System (DNS) and the Dynamic Host Configuration Protocol (DHCP) are two examples of protocols that use both TCP and UDP for different functions and can use the same port numbers for both (as DNS does), or different port numbers depending on whether the protocol is using TCP or UDP as its transport mechanism (as DHCP does).

There are also proprietary protocols that are monolithic in nature, and a single protocol may span multiple layers or functions in the protocol stack. Other examples are protocols that encapsulate other protocols, such as the Layer 2 Transport Protocol (L2TP), which, when used in VPNs, encapsulates IPSec, the security protocol that protects the data.

Regardless of the layer at which multilayer protocols function, the key is that you must secure these protocols based on several criteria and consider how security issues affect

the protocols used across different layers. Considerations in securing multilayer protocols include

- The layers in which the protocol resides
- The port(s) the protocol uses, as well as whether its transport protocol is TCP or UDP
- The devices at the layers in question, such as a router or switch
- Whether or not encapsulation or encryption are used
- Whether secure protocols are used to protect nonsecure protocols, such as the use of IPSec to protect unsecured VPN traffic, or TLS used to protect HTTP traffic

Converged Protocols

Over the history of networking there have been other types of traffic, such as voice, for instance, that have used separate equipment, routes, methods, and even protocols to get the information in various forms from one point to another. Slowly, these different technologies have converged to all use standardized networks and data structures. Most of this information now can be carried over a standard TCP/IP network as digital data, in fact. The CISSP exam objectives require you to understand a few of these *converged protocols*, but note that there are many more that we don't address here. The converged protocols specifically listed in exam objective 4.1 are Fibre Channel over Ethernet, iSCSI, and Voice over Internet Protocol traffic.

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) is a technology that allows storage area network (SAN) solutions to function over Ethernet. FCoE operates at high speeds, using the Fibre Channel protocol, and requires specialized network equipment, such as switches, host-based adapters, network interface cards, and special cabling. Fibre Channel is a serial data transfer protocol used for high-speed data storage and transfer over optical networks. It has been updated to allow its use over copper Ethernet cables, hence FCoE. Using FCoE, SANs can be built with extremely low latency, allowing the actual storage location to be unimportant and transparent to the user, regardless of whether it is on the same local network or half the globe away over WAN links. FCoE works at layer 2 of the OSI model and requires a minimum speed of around 10 gigabits per second (Gbps) to work. A follow-on technology is Fibre Channel over IP (FCIP), which doesn't require any specific network speed or devices, as it is encapsulated within IP and functions at layer 3 of the OSI model, the network layer.

Internet Small Computer Systems Interface

Like FCIP, the Internet Small Computer Systems Interface (iSCSI) is a storage technology that works at layer 3 of the OSI model, as it is based on IP. It too can be used for location-agnostic data storage and transmission over a LAN or over long-distance WAN links. iSCSI runs over standard network equipment, unlike FCoE, which requires specialized equipment, so iSCSI is not as expensive as FCoE.

iSCSI essentially sends block-level access commands to network-based storage devices; in other words, it sends SCSI data transfer commands over a TCP/IP network. Clients, called *initiators*, send these commands to storage devices, called *targets*, on remote networks for data storage and retrieval.

Voice over Internet Protocol

Until the past decade or so, voice was considered a separate type of analog data and was transmitted over separate, unique systems. However, voice has been integrated into digital data networking as speed and efficiency of network equipment and protocols have matured. We will go into more depth on Voice over Internet Protocol (VoIP) and its transition to becoming a part of network data traffic in Objective 4.3.

Cross-Reference

VoIP is discussed in detail in Objective 4.3.

Micro-segmentation

Networks are segmented for a variety reasons, mostly for performance and security. For performance, we segment networks using devices such as routers and switches to eliminate broadcast and collision domains, which in turn reduces the amount of congestion on networks. For security, we segment networks so that sensitive hosts and networks are separated from the general population. For example, the organization's servers and network devices can be logically separated on their own virtual LAN (VLAN), rather than just physically segmented away from the general user population, so that rules and filters can be applied to control access to those devices and networks.

Micro-segmentation is an extension of this strategy, sometimes segmenting all the way down to a sensitive host or even an application so that it is separate from the general network and can be properly secured. There are several different ways we can segment networks and hosts at the micro level, including putting them on their own logical subnet, using VLANs, complete encryption of all traffic to and from a host or segment, and other logical methods. We can also segment physically using different cabling and media, or even completely disconnecting hosts from the network altogether, forcing the use of manual methods (e.g., sneaker net) to transfer data to and from those isolated hosts. Virtualization is one of the logical methods we can use to segment networks and hosts, and we will discuss several methods of network virtualization next.

Software-Defined Networking

Just as operating systems and applications can reside in a virtualized environment, so can networks. Software-defined networking (SDN) is a virtualization technology that allows for traditional networks to be virtualized using software to control how traffic is forwarded between hosts within the organization. This type of virtualization allows software to control routing

tables and decisions, bandwidth utilization, quality of service, and so on. SDN uses a software controller to handle dynamic traffic routing, which eliminates some of the slower, hardware infrastructure–based decisions involved in physical networks. SDN also allows quick network reconfiguration and provisioning.

SDN takes over functions of the *control plane* (part of the control layer of the infrastructure, which assists in updating routes) and uses software-based protocols to decide how and where to send traffic. SDN separates these functions from network hardware, which can be slower and less responsive to dynamic or manual changes to the network topology. SDN off-loads work from the forwarding plane to reduce the complex logic that goes into those functions. The *forwarding plane* is the part of the infrastructure that actually makes traffic forwarding decisions, based on the routing information from the control plane, and is usually implemented as a hardware chip in network devices. SDN separates the control and forwarding planes and makes those decisions on behalf of the hardware.

SDN offers organizations a greater flexibility in level of control over traffic within a network. Organizations are no longer required to use only a specific vendor whose products are only interoperable with each other. SDN uses open standards and protocols and is vendor agnostic.

Software-Defined Wide Area Networking

Software-defined wide area networking (SD-WAN) is the logical extension of SDN and can be used to virtualize connectivity and traffic control between geographically separated infrastructures, such as those that are found in remote locations, cloud services, and multiple data centers. SD-WAN, like SDN, can control traffic between sites, providing bandwidth and QoS control.

Virtual eXtensible Local Area Network

Virtual eXtensible Local Area Network (VxLAN) is a protocol (defined in RFC 7348) that encapsulates VLAN management traffic and allows it to be sent across logical and physical subnets to geographically separated locations. This expands VLANs beyond locally created virtual subnets and allows VLANs to extend to different geographic locations separated by WAN links. VLANs are normally limited to a single router interface, but VLANs that use VxLAN encapsulate the layer 2 VLAN frames into UDP so they can be sent via IP over links that cross multiple subnets. VLANs that use VxLAN are also far more scalable; VLANs are normally limited to 4096 networks (due to the 12-bit VID), but VxLAN expands that to over 16 million virtualized networks.

Encapsulation

Encapsulation was introduced earlier when we discussed how data is wrapped within other data as it moves down the OSI model layers. The data from the layer above is repackaged into another PDU, with header information added, and that entire package becomes the data for the next layer down in the stack. Encapsulation is much more than that, however, and is very powerful in protecting network data and traffic. Encapsulation can also be used to segment entire networks.

Consider a network where most of the data is not sensitive, but certain sensitive data is transmitted using encryption protocols, essentially segregating that data from the rest of the network. The sensitive data can only be sent and received by specific hosts that have the ability to encrypt and decrypt it. Encapsulation, or tunneling, is also what segments VPN traffic from the larger public Internet. A VPN connection tunnels sensitive traffic through a nonsecure network, such as the Internet, where it is received by a remote VPN concentrator at the destination network and de-encapsulated and decrypted. These are all examples of how encapsulation (as well as tunneling, which is a form of encapsulation) works to protect data by segmenting the network traffic away from untrusted networks.

Cross-Reference

We will also briefly discuss virtualized networking technologies, such as VLANs, SDN, VxLANs, and SD-WANs, in Objective 4.3.

Wireless Technologies

Networking uses two types of media: wired and wireless. We'll discuss securing wired media in Objective 4.2 (and briefly mention wireless media as well) and focus on wireless networking in this section. Keep in mind that wireless includes the use of a wide range of technologies in addition to Wi-Fi, including radio frequency (RF) signals, microwave and satellite transmissions, infrared, and other technologies. We'll also discuss cellular technologies. This section will not teach you everything you need to know about wireless technologies; this is simply a review that focuses on the critical technical and security aspects of wireless networking that you need to know for the CISSP exam.

Wireless Theory and Signaling

We can't discuss wireless networking without talking about a few of the standards and technologies on which wireless networking is based. Wireless uses the electromagnetic spectrum, which consists of a variety of energy types, including radio, microwave, and light. We will touch on each of these energy types, but first we must discuss a little bit of basic wireless theory that you should keep in mind for the exam.

The electromagnetic (EM) spectrum can be divided into many types of energy, but for exam purposes we will discuss primarily the radio frequency (RF) portion of the spectrum; keep in mind that the concepts apply to many other areas of the EM spectrum as well. The fundamental RF signal is a *sine wave*, which appears as an electrical current changing voltage uniformly over a defined period of time. The wave changes direction from one side to the other as time passes. As the wave changes direction over a measured period of time, it is known as a *cycle*. The measure of cycles per second is known as its *frequency*. Frequencies are measured by a unit called hertz (Hz), which use prefixes such as kilohertz (KHz, or 1000 Hz), megahertz (MHz, or 1 million Hz), and gigahertz (GHz, which is 1 billion Hz).

Frequencies within a specific range are called the *bandwidth*, which is the difference between the upper and lower frequency in the range. Radio waves can be changed, or *modulated*, based on *amplitude* (height of the sine wave) or its frequency. Because radio waves normally exhibit a predictable pattern, they should all be uniform given the same frequency and modulation. However, when they are not uniform, they are said to be out of *phase*, which can result in a garbled transmission. RF propagation suffers from many issues, including *absorption* (absorbed by material), *refraction* (resulting in the waves being bent by an object), *reflection* (when waves are reflected off an object), *attenuation* (gradual signal weakening over time and distance), and *interference* (noise interfering with the construction of the waves, resulting in an incomplete or garbled transmission).

Antennas are used to send and receive radio signals and are essentially conductors of RF energy. Antennas come in many different shapes and sizes, depending on the nature and characteristics of the RF signal being used. Radio signal power is measured in watts (W), although typically measurements such as milliwatts are used to describe the lower-power transmitters, such as those used in wireless networking.

Wireless networking uses several different signaling methods. Many of these involve the use of *spread-spectrum* technologies, in which individual signals are spread across the entire *frequency band*, or section of allocated frequencies. This allows a transmitter to use bandwidth more effectively, since the transmitting system can use more than one frequency at a time. These signaling methods include the following:

- **Direct sequence spread spectrum (DSSS)** Uses the entire frequency band continuously by attaching chipping codes to distinguish transmissions; both the sender and receiver must have the correct chipping codes to communicate.
- **Frequency hopping spread spectrum (FHSS)** Uses portions of the bandwidth in a frequency band, splitting it into smaller subchannels, which the transmitter and receiver use for a specified amount of time before moving to another frequency or channel.
- **Orthogonal frequency division multiplexing (OFDM)** This is not a spread-spectrum technology, but it is an important signaling method used in modern wireless networking. OFDM is a digital modulation scheme that groups multiple modulated carriers together, reducing bandwidth. Modulated signals are *orthogonal*, or perpendicular, to each other, so they do not interfere with other signals.



EXAM TIP The key points about RF theory to remember concerning wireless networking are frequency and the signaling methods used.

Next, we will discuss wired standards that use the various radio frequency bands and signaling methods.

Wi-Fi

Wi-Fi is the name given to a variety of technologies used to allow homes and businesses to make use of the electromagnetic spectrum to connect devices to each other and the Internet to send and receive data, foregoing wired connections whenever possible, and creating the "mobile revolution." Note that the term Wi-Fi is a term trademarked by the Wi-Fi Alliance. Now that we have briefly touched on the science behind the technology, we will discuss Wi-Fi technologies, including the fundamentals and various important standards used, and then discuss Wi-Fi security.

Wireless standards to focus on for the CISSP exam begin with the IEEE 802.11 standard and its amendments. These are the standards assigned to wireless networking, and there are many other standards that also contribute to wireless networking. Some of these standards dictate frequency usage and signaling method, while others dictate quality of service and security. We will not cover every single wireless standard that exists, but there are several you should be familiar with for the exam.

Wi-Fi Fundamentals

Wireless LANs (WLANs) use devices that have radio transmitters and receivers installed in them, such as smartphones, laptops, PCs, tablets, and so on. While these devices can directly communicate with each other over Wi-Fi (this is called an *ad hoc* network, and can be both problematic to set up and unsecure), most home and business wireless networks use wireless access points (WAPs). A WAP manages the wireless connection between the access point and the client device, as well as between devices. When using a WAP, this is known as *infrastructure mode*, and devices attached to the WAP are referred to as a Basic Service Set (BSS). The WAP may also be connected to a wired network so that wireless clients can access a larger network. To use a WAP, clients must be configured with the correct *Service Set Identifier (SSID)*, which is essentially the wireless network name configured on the WAP. The SSID may be visible to other wireless clients that are not part of the network, or it may be hidden by not broadcasting it so that unauthorized wireless clients cannot easily join the network. Note that wireless devices must be on the same frequency band (channel) to connect with each other, as well as share common security parameters, which we will discuss a bit later. Frequency bands for the 802.11 standards for wireless networking include those in the Industrial, Scientific, and Medical (ISM) ranges (900 MHz, 2.4 GHz, and 5.8 GHz, respectively), and the Unlicensed National Information Infrastructure (UNII) band (5.725 GHz to 5.875 GHz, which overlaps to a small degree with the upper ISM band).

Wi-Fi Physical Standards

Wi-Fi physical standards were developed over many years and are primarily those published by the Institute of Electrical and Electronics Engineers (IEEE), along with the same standards that are established for wired networks, such as Ethernet, Fiber Distributed Data Interface (FDDI), and Token Ring. Most Wi-Fi standards fall under the IEEE 802.11 standard and its amendments.

There are amendments for data rates, physical (PHY) signaling technology, frequency band, and range. There are also amendments for security, quality of service, and other characteristics that wireless networks and associated equipment must meet. The Wi-Fi Alliance, a consortium of wireless equipment manufacturers and professional bodies, has contributed to many of the standards and promulgates them as the official standards of the wireless industry. Devices sold by manufacturers must comply with Wi-Fi Alliance standards to be certified by the body.

Table 4.1-2 summarizes some of the more important IEEE 802.11 amendments that you should be familiar with for the CISSP exam.

Wi-Fi Security

Wireless security has been problematic since the original 802.11 specifications were released. The first attempt to bring security to wireless transmissions was through Wired Equivalent Privacy (WEP), prevalent in earlier wireless devices, such as those that connected to 802.11a and 802.11b networks. However, WEP used weak initialization vectors and a problematic implementation of the RC4 streaming cipher. WEP was considered highly unsecure and easily breakable, so it has long since been deprecated. Later, Wi-Fi Protected Access (WPA) was developed by the Wi-Fi Alliance while waiting on a formal IEEE standard to be implemented.

WPA allowed larger key sizes and implemented the Temporal Key Integrity Protocol (TKIP). WPA was implemented on many wireless devices manufactured before the formal IEEE security standard, 802.11i, was implemented, so many devices during that time were not only backward compatible with WEP but also compatible with the interim WPA standards and the newer official WPA2 standards, as the IEEE 802.11i standard came to be known. WEP has been deprecated because it can be rapidly cracked, as has been WPA. WPA2's improvements include the ability to use the Advanced Encryption Standard (AES) over TKIP, as well as larger key sizes and better encryption methods.

Many modern devices still use WPA2, but most are transitioning to the new WPA3 standard, required by the Wi-Fi Alliance since 2020. While not an official IEEE standard, the Wi-Fi Alliance has ensured that WPA3 continues to implement the requirements of the original 802.11i amendment, as well as IEEE 802.1s (introducing Simultaneous Authentication of Equals [SAE] exchange, which replaces the need for preshared keys used previously in Wi-Fi security protocols), and IEEE 802.11w, which provides for protection of management frames. Implementation of WPA3 is mandatory on all devices certified by the Wi-Fi Alliance after July 2020.

In addition to running recommended security protocols, such as WPA3, the following are several other measures you should take to secure wireless networks:

- Use deprecated security protocols only when absolutely necessary or when unable to upgrade or replace equipment, and mitigate weaknesses with other controls (e.g., IPSec).
- Ensure passwords and other authenticators maintain a minimum character length and complexity.
- Ensure wireless access points and network equipment are physically protected.
- Reduce transmitting power of wireless access points and other devices to only what is necessarily for effective coverage.

TABLE 4.1-2 Common Wireless Networking Standards

802.11 Standard or Amendment	Generation	Data Rates	PHY Signaling Technology	Band	Frequency Range
802.11	Wi-Fi 0 (1997)	1 to 2 Mbps	FHSS/DSSS	ISM, 2.4 GHz	2.4 GHz
802.11b	Wi-Fi 1 (1999)	1–11 Mbps	DSSS	ISM, 2.4 GHz	2.4–2.4835 GHz
802.11a	Wi-Fi 2 (1999)	6–54 Mbps	OFDM	UNII, 5 GHz	5.150–5.250 GHz UNII-1 5.250–5.350 GHz UNII-2 5.725–5.825 GHz UNII-3
802.11g	Wi-Fi 3 (2003)	Up to 54 Mbps (when used in mixed mode, data rates match legacy devices)	OFDM (but backward compatible with 802.11b using DSSS and HR/DSSS)	ISM, 2.4 GHz	2.4–2.4835 GHz
802.11n	Wi-Fi 4 (2008)	Up to 600 Mbps (when used in mixed mode, data rates match legacy devices)	HT-OFDM	ISM, 2.4 GHz UNII, 5 GHz	Same as 802.11a/b/g
802.11ac	Wi-Fi 5 (2014)	433–6933 Mbps	VHT-OFDM	UNII, 5 GHz	Same as 802.11a/n
802.11ax	Wi-Fi 6/6e (2019/2020)	600–9608 Mbps	HE-OFDMA	ISM, 2.4 GHz UNII, 5 GHz, 6 GHz	Same as 802.11a/b/g/n and 5.925–7.125 GHz

- Don't rely on hiding or simply not broadcasting the network's SSID to provide any security, since this deters only casual snoopers, not experienced hackers.
- Periodically change authenticators used over wireless networks in accordance with security policies.
- Use enterprise-level mutual authentication for both users and devices when connecting to the corporate network.
- Use port-based authentication (802.1X) as part of enterprise-level security.
- Don't rely on MAC address filtering to provide any security against rogue clients connecting to the WAP, since MAC addresses can be easily spoofed.
- Actively scan the network for rogue devices, such as unknown WAPs, which may be used to conduct an evil twin attack.

Bluetooth

Bluetooth, an IEEE 802.15 standard, helps to create wireless personal area networks (WPANs) between devices, such as between a smartphone and headphones or external speakers, between a computer and keyboard and mouse, and between a myriad of other Bluetooth-capable devices. This connection process is called *pairing*. Bluetooth uses some of the same frequencies used by 802.11 devices, in the 2.4-GHz range. Bluetooth has a maximum range of approximately 100 meters, depending on the versions in use, environmental factors, and the transmission strength of the transmitting device. At the time of this writing, Bluetooth is currently in version 5.3.

Earlier versions of Bluetooth were susceptible to attack by receiving unsolicited pairing requests from devices, resulting in unrequested messages being sent to the receiver, such as ads, harassing messages, and even illicit material such as pornography. This type of attack is called *Bluejacking*. Another type of attack is called *Bluesnarfing*, which is more invasive and allows an attacker to request a Bluetooth connection, and, once paired, to read, modify, or delete information from the victim's device. Both of these attacks can be prevented by simply making the device nondiscoverable unless the owner is intentionally pairing it with a known device, and also changing the default factory code (or PIN), which is often required to pair Bluetooth devices. The default PIN might be something simplistic, such as 0000, for example, and may be commonly used on many different devices.

Bluetooth Low Energy (BLE) is a version of Bluetooth designed for use in devices that require low power, such as medical devices and Internet of Things (IoT) devices. Note that BLE is not compatible with standard Bluetooth, although both types may be found on the same device.

Zigbee

Zigbee is a technology based on the IEEE 802.15.4 standard, with very low power requirements and a correspondingly low data throughput rate. It requires devices to be in very close proximity to each other and is frequently used to create WPANs. Zigbee is used in IoT applications and is able to provide 128-bit encryption services to protect its transmissions. It is used in industrial control systems, medical devices, sensor networks, and even home automation, such as the type used to control lights or temperature in a smart home.

Although Zigbee has encryption services built-in, it uses what is known as an *open trust* model, meaning that all applications and devices on a Zigbee network inherently trust each other. However, network management and data protection are secured using three different 128-bit symmetric keys. The *network* key is the one shared by all devices on the Zigbee network and is used for broadcast messages. A *link* key is used for each pair of connected devices and is used during unicasts between two different devices. The *master* key is unique for each connected pair of devices to establish symmetric keys and key exchange. Although Zigbee has encryption capabilities built into the standard, the open trust model allows almost any other device to authenticate to it, so the protocol's security mechanisms are lacking. Zigbee depends on a well-protected physical environment to ensure its devices are secured.

Satellite

Satellites can be used to provide wireless network access as a link between two distant points. Satellites are a line-of-sight technology, meaning that all points sending and receiving data using the satellite must be within the satellite's direct line of sight. The *footprint* of the satellite is the area of coverage. Most satellite wireless clients (called ground stations) communicate with a centralized hub over normal wired or wireless links (including terrestrial microwave and normal Wi-Fi), but even remote end-user stations have the capability of reaching the satellite for transmitting and receiving signals. A transponder on the satellite is used to transmit signals to and receive signals from the ground, and normally the antenna used to receive and transmit satellite signals on the ground is the shape of a dish. Satellites can be used for both broadband television and Internet access.

Satellites normally are in one of two types of orbits so that they can provide communications services. Low Earth orbit (LEO) satellites operate at a distance between 99 and approximately 1,245 miles above the Earth's surface, so there is not as much distance between ground stations and satellites. The shorter distance means that smaller, less powerful receivers can be used. However, this also means that there is less bandwidth available. LEO satellites are frequently used for international cellular communications and by satellite Internet providers.

Geosynchronous satellites orbit at an altitude of 22,236 miles and rotate at the same rate as the Earth. This has the effect of making satellites appear fixed in orbit over the same spot. A geosynchronous satellite requires a large ground dish antenna, however, and, because of the distance, introduces a lot more latency into the communications process than LEO satellites do. Geosynchronous satellites provide services for transatlantic communications, TV broadcast services, and so on.

Li-Fi

Li-Fi uses a light to transmit and receive wireless signals, since it is also an electromagnetic wave with a much higher frequency. Theoretically, light can carry much more information than normal radio waves, as evidenced by fiber-optic cables that are used for high-throughput backbones. Li-Fi is constrained to a much smaller space than normal RF wireless using the 802.11 standard frequencies, so it is much harder to intercept. Li-Fi is only in the beginning stages, so the technology is not yet mature; however, it has the capacity to support much more data with a

lower latency and better adaptability in locations where lower frequency RF waves may be prone to interference.

Cellular Networks

Cellular technology is a ubiquitous form of wireless media with the proliferation of so many mobile devices, particularly smartphones and tablets. Cellular networks get their name from being part of a geographical area known as a "cell," which is the maximum transmission and receiving distance of a cellular tower and its associated base station. Cells are hexagonal in shape and located adjacent to each other, so cellular technologies provide for seamless handoff between a mobile device and a cell tower as the device moves from one cell to another.

Because only a finite number of frequencies are allocated to a cellular network and a given cellular tower and base station, devices must contend for the use of those frequencies. To address that issue, many technologies have been developed that allow for *multiple access* to these frequencies. Usually, multiple access means using different techniques such as time division, code division, or frequency division to allow multiple devices to access the frequency band. Cellular technologies use the following multiple access methods:

- **Time division multiple access (TDMA)** Signaling method that uses time slices of a particular frequency, allowing each user to use the frequency for a short, finite period of time.
- **Code division multiple access (CDMA)** Spread-spectrum signaling method that allows multiple users to use a frequency range by assigning unique codes to each user's data transmission and device.
- **Frequency division multiple access (FDMA)** In this method, the available frequency range is divided into sub-bands, or *channels*, and each channel is assigned to a particular subscriber (device) for the duration of the subscriber's call; the subscriber has exclusive use of this channel for the duration of the session.
- **Orthogonal frequency division multiple access (OFDMA)** This is similar to FDMA except that the channels are subdivided into closely spaced orthogonal frequencies with narrow bandwidths; each of these different subchannels can be transmitted and received simultaneously.



EXAM TIP Don't confuse OFDM (orthogonal frequency division *multiplexing*) with OFDMA (orthogonal frequency division *multiple access*). OFDM is a single-user digital multiplexing scheme and can send multiple types of signals over multiple carriers, whereas OFDMA is an extension of OFDM implemented for multiple users to share cellular frequencies, with three times higher throughput than OFDM.

Cellular services are often referred to by their generation, or "G." Generations of similar technologies include 2G, 3G, 4G, and 5G. Modern mobile devices rely on at least a minimum of 4G technologies to adequately send and receive voice and multimedia, which consists of

TABLE 4.1-3 Summary of Cellular Generations and Characteristics

Generation (G)	Signaling Characteristics	Description
1G	Analog transmission of voice for circuit-switched networks	Voice only, no data
2G (1993) 2.5G	TDMA, CMDA; 1800 MHz spectrum, 25 MHz bandwidth	First cellular generation to send both voice, and limited low-speed circuit-switched data Higher bandwidth and "always on" technology
3G (2001)	FDMA, TDMA, CDMA; 2 GHz spectrum, 25 MHz bandwidth	Integration of voice and data; transition to packet-switching over circuit-switching technologies
3.5G (3GPP) 4G/LTE (2009) 5G (2018)	OFDMA 100 MHz bandwidth (4G) 3 GHz to 86 GHz spectrum and 30 MHz to 300 MHz bandwidth (5G)	Higher data rates All IP packet-switched networking Higher frequency ranges but prone to more interference; data rates of up to 20 Gbps

video, text, audio, and Internet content. 4G was the first stable technology to primarily use IP to send and receive data. Transmissions from cellular devices are normally encrypted between the device and the cell tower, but once the tower receives them, they are transmitted over the normal long-haul wired telephone infrastructure, so those transmissions may be unprotected and sent in clear text, making them susceptible to interception. Cellular technologies and their generations include those listed in Table 4.1-3.



EXAM TIP While you may not see specific questions regarding legacy cellular technologies (1G through 3G) on the exam, it's still a good idea to familiarize yourself with them to understand how they lead into modern technologies such as 4G and 5G.

Content Distribution Networks

A *content distribution network* (or *CDN*, also called a *content delivery network*) is a multilayered group of network services and resources, which are implemented at various levels all over the entire Internet, including in data centers and over WAN links. Large multimedia providers, such as streaming media companies (e.g., Hulu, Netflix, Amazon, and so on), use CDNs to provide high content availability and performance while keeping latency to a minimum. CDNs provide redundancy and load balancing services to these providers, ensuring high data availability to their customers. Related to CDNs is the concept of *edge computing*, which also helps

reduce latency and provides high availability by locating equipment geographically closer to users. Implementing CDNs reduces the geographic distance that content has to travel to and from the user, which lowers latency by eliminating the need to download content from halfway across the world over multiple WAN links.

Cross-Reference

Edge computing was discussed in more detail in Objective 3.5.

REVIEW

Objective 4.1: Assess and implement secure design principles in network architectures In this objective we briefly discussed the highlights of networking concepts and fundamentals, such as the OSI model and TCP/IP stack. We also discussed IP networking and the basics of secure protocols. These key elements are necessary to ensure that secure design principles are adhered to when designing networks. We examined the application of secure networking concepts and surveyed multilayer protocols and converged technologies. We also touched on the key concepts of micro-segmentation, software-defined networking, and Virtual eXtensible Local Area Network. We explored wireless and cellular networks and discussed security concerns with those technologies. Finally, we discussed the purpose of content distribution networks in terms of providing better availability to multimedia consumers.

4.1 QUESTIONS

1. You have configured IPSec on your local network, and sensitive traffic is being sent between specific hosts. However, when you look at the traffic in a protocol analyzer, it is not encrypted. Which of the following protocols must you ensure is configured correctly for IPSec to encrypt traffic?
   - **A.** Encapsulating Security Payload (ESP)
   - **B.** Authentication Header (AH)
   - **C.** Internet Key Exchange (IKE)
   - **D.** Internet Security Association and Key Management Protocol (ISAKMP)
2. Your organization makes extensive use of VLANs. Because it has merged with another company, it now has regional offices all over the globe and has also moved some of its infrastructure into the cloud. You wish to retain the ability to use VLANs that can be deployed regardless of geographic location or WAN links. Which of the following technologies should you recommend that your company implement?
   - **A.** Software-defined networking (SDN)
   - **B.** Software-defined wide area networking (SD-WAN)
   - **C.** Virtual eXtensible Local Area Network (VxLAN)
   - **D.** Fibre Channel over Ethernet (FCoE)

3. Which of the following IEEE 802.11 standard amendments specifies Wi-Fi operation in the 6-GHz band?
   - **A.** IEEE 802.11ac
   - **B.** IEEE 802.11g
   - **C.** IEEE 802.11i
   - **D.** IEEE 802.11ax
4. You are teaching a class on wireless technologies to your company's IT staff. One student asks you about the different multiple access methods available to help manage frequency band usage in cellular networks. Which of the following correctly describes the three methods for multiple access used?
   - **A.** Frequency division, code division, and bandwidth division
   - **B.** Time division, amplitude division, and frequency division
   - **C.** Time division, code division, and frequency division
   - **D.** Time division, bandwidth division, and frequency division

4.1 ANSWERS

1. **A** Encapsulating Security Payload (ESP) is the protocol in IPSec that is responsible for encrypting data.
2. **C** Virtual eXtensible Local Area Network (VxLAN) is a protocol that encapsulates VLAN management traffic and allows VLAN technology to be extended over WAN links across multiple geographic locations.
3. **D** Of the choices given, the only IEEE 802.11 standard that operates in the 6-GHz band is IEEE 802.11ax.
4. **C** The three primary methods for managing frequency usage in cellular networks are time division multiple access (TDMA), code division multiple access (CDMA), and frequency division multiple access (FDMA). Neither bandwidth nor amplitude is used to manage multiple device access to frequencies.



Objective 4.2 Secure network components

In this objective we will discuss the minimal security controls all infrastructure components should be required to maintain. While this is only a brief review, the key takeaways for this objective are that network components should be secured both physically and logically from unauthorized access, and that security controls, such as strong authentication, encryption, configuration management, and anti-malware, are a must.

Network Security Design and Components

Network components include not only network infrastructure devices and servers, but also network hosts, such as workstations and mobile devices, and transmission media. Securing these components is discussed throughout this book, but this objective specifically points out controls that you must consider as the minimum necessary to ensure secure network operation and asset protection.

Operation of Hardware

Infrastructure hardware includes servers, switches, and routers, in addition to security devices and any other network-enabled hardware that provides essential services to users. Infrastructure hardware encompasses not only traditional IT devices but also nontraditional devices identified as industrial control systems and Internet of Things (IoT) devices. Each of these types of devices has its own strengths and weaknesses, along with its own configuration management processes. Network devices allow varying degrees of control over security configuration and the traffic that flows through them. They also can each be secured to different degrees; cutting-edge technology is likely to have better security controls built-in than legacy systems. You must consider many aspects of securing network devices, including network architecture, network traffic control devices, and specialized security devices.

Network Architecture

Network architecture design is the first step in securing network components. You should thoroughly consider placement of devices and how they are connected as a security control. Network architectures can be very simplistic or they can be very complex and include physical and logical segmentation (e.g., VLANs, different IP subnetworks, host isolation, perimeter devices, etc.). In addition to providing security through isolating and protecting sensitive hosts, network architectures can also reduce network traffic issues, such as congestion that results from collisions and broadcasts. With that in mind, there are various network architectures you should understand for the CISSP exam:

- **Intranet** A private network residing only within an organization and separated from the public Internet or other nontrusted networks
- **Extranet** Specially segmented portion of an organization's network that is configured to provide services for business partners, customers, and suppliers
- **Demilitarized zone (DMZ)** Carefully controlled perimeter network that is used as a barrier between public or nontrusted networks and internal or protected networks
- **Virtual LAN (VLAN)** Not a network architecture per se, but a method of using advanced switches to create virtual (rather than physical) local area networks with their own IP address space—essentially, software-defined LANs that can be logically separated from other network segments and have their own access control rules
- **Bastion host** Specially configured (hardened) host that separates untrusted networks or hosts from sensitive ones

Note that these architectures could consist of many different devices arranged in specific configurations to provide protection and control traffic flow, or even a single device, like the bastion host mentioned, that separates sensitive networks from untrusted ones.

Network Hardware Devices

Network hardware devices serve many different functions, including moving traffic from one network or node to another using physical or logical addressing; ensuring that traffic gets delivered to the correct segment or host; and filtering traffic based on its characteristics, such as port, protocol, service, address, and so on. Network hardware devices also work to raise the quality of service and reliable delivery for specific types of traffic or content. Although this section is not a comprehensive review of network hardware devices, the most common devices you will encounter are presented in Table 4.2-1 as a refresher.

TABLE 4.2-1 Common Network Hardware Devices

Device	Description	Security Features
Repeater/ concentrator/ amplifier	Amplifies and regenerates network signals	None
Hub	Simple network connection point for devices	None
Modem	Modulator–demodulator that carries digital data over analog voice lines	None
Bridge	Connects two or more network segments together using same layer 2 standards (e.g., Ethernet)	None
Switch	Connects two or more network segments together; can break up collision domains (or broadcast domains when implementing VLANs)	Hardware address filtering, sniffing prevention, VLAN
Router	Routes traffic between LANs based on logical IP addressing	Logical address and domain filtering, port/ protocol/service filtering
Brouter	Combines bridge and router functionality; separates both collision and broadcast domains; functions as router but will attempt to bridge if routing fails	Physical and logical address and domain filtering, port/protocol/ service filtering
Gateway	Connects disparate networks that use different protocols or topologies	Application-level filtering
Proxy	Separates trusted from untrusted networks; makes requests on behalf of hosts	Application content and protocol filtering; stateful inspection

Firewalls

Firewalls are ubiquitous network security devices; they are the foundation of establishing a secure perimeter protecting an organization's enterprise network. Firewalls are used to control and regulate traffic between networks, such as the Internet and the internal network, between DMZ networks and external networks or extranets, and even between sensitive hosts and segments inside the private network.

Firewalls are used to filter (by blocking or allowing) traffic based on different criteria, specified by elements in its rule sets. Elements of firewall rules include basic characteristics of traffic, such as source or destination IP address, domain, port, protocol, and service, but also more complex traffic characteristics, such as user context, time of day, anomalous traffic patterns, and so on. Many next-generation firewalls (NGFWs) are combination devices that include features and characteristics of proxies, intrusion detection/prevention devices, and network access control.

Firewalls usually have more than one interface, which allows them to connect to and filter traffic between multiple networks. Firewalls can also be deployed in multitier architectures, such as those that might be found in a demilitarized zone. For example, you could have a single-tier firewall that separates the Internet from an internal network, or a two- or three-tier firewall setup that divides the network into protective zones, each more restrictive.

From both a historical and functional perspective, firewalls can be classified in terms of type and level of functionality. The different types of firewalls include

- *Packet-filtering* or *static* firewalls filter based on very basic traffic characteristics, such as IP address, port, or protocol. These firewalls operate primarily at the network layer of the OSI model (TCP/IP Internet layer) and are also known as *screening routers*; these are considered first-generation firewalls.
- *Circuit-level* firewalls filter session layer traffic based on the end-to-end communication sessions rather than traffic content.
- *Application-layer* firewalls, also called *proxy* firewalls, filter traffic based on characteristics of applications, such as e-mail, web traffic, and so on. These firewalls are considered second-generation firewalls, which work at the application layer of the OSI model.
- *Stateful inspection* firewalls, considered third-generation firewalls, are dynamic in nature; they filter based on the connection state of the inbound and outbound network traffic. They are based on determining the state of established connections. Stateful inspection firewalls work at layers 3 and 4 of the OSI model (network and transport, respectively).
- *Next-generation* firewalls (NGFWs) are typically multifunction devices that incorporate firewall, proxy, and intrusion detection/prevention services. They filter traffic based on any combination of all the other firewall techniques, to include deep packet inspection (DPI), connection state, and basic TCP/IP characteristics. NGFWs can work at multiple layers of the OSI model, but primarily function at layer 7, the application layer.



EXAM TIP You should understand the various types of firewalls and their functions, as well as at which layers of the OSI model they function.

Network Access Control Devices

Network access control (NAC) is a method of strictly controlling access to the organization's infrastructure by requiring stringent security configurations for devices. The goal of NAC is to enforce strong authentication and security policies to ensure only hosts meeting minimum security requirements can connect to the network. NAC is often implemented through devices or dedicated appliances but can also be implemented through software on a traditional security device, such as an IDS/IPS or next-generation firewall.

NAC devices are physically protected similarly to other infrastructure equipment; they should be kept in a secure locked room with limited access. Administrative access to NAC devices should be very limited and should be controlled through strong authentication protocols. You should also restrict access to the devices that can successfully connect to the infrastructure through the NAC devices; policy should determine whether users can bring their own devices and connect to the corporate network, for example, and if this is not the case, only a predefined list of authorized devices should be able to connect through the NAC devices and into the corporate network.

Since NAC devices are used to initially grant or deny access to the network from authorized devices, make sure that each authorized device is required to undergo a health or compliance check, to ensure that it has the appropriate patches installed, configuration settings secured, and up-to-date anti-malware. NACs should also be monitored and audited on a continual basis.

Network Device Security Controls

In general, technical security controls used to protect network devices are common across the infrastructure, but you should make sure that you address the following technical controls for critical network devices:

- Limit both physical and remote access to only authorized administrators.
- Use only secure protocols for remote administration, such as Secure Shell (SSH), a secure VPN protocol such as IPSec, or Transport Layer Security (TLS); use multifactor authentication whenever possible.
- Configure complex passwords for each device.
- Use role-based administration for network devices.
- Turn off all unnecessary services, especially unsecure services such as Telnet.
- Audit all administrative access to each device.
- Limit which hosts administrators can use to access network devices.
- Ensure all devices are updated with the latest security patches.

Physical controls are also very important for network device security. Obviously, network devices should be locked in a secure area, such as a data center, communications closet, or server room. This secure processing area should have very limited personnel access, and any uncleared personnel, such as physical maintenance personnel, for instance, should be escorted when in the area around network devices. You should also restrict any type of personal device that transmits using wireless or cellular signals when in the vicinity of network devices, such as smartphones, tablets, and laptops.

In addition to technical and physical controls, you should also consider controls that can directly impact the availability goal of security:

- Ensure temperature and humidity levels are properly maintained in locations where network devices operate.
- Use both redundant power supplies and backup power devices, such as uninterruptible power supplies and generators, for network devices.
- Be aware of end-of-life support for network devices and maintain the warranties on those devices.
- Ensure that spare parts and components are available for all network devices.
- Maintain an accurate and up-to-date inventory of all authorized network devices.

Transmission Media

Transmission media can be categorized as either wired media or wireless media. Each category has some common security requirements you should pay attention to, but each category also has its own unique requirements. For wired media, the key is to protect cabling from unauthorized physical access. You should secure cabling in the following ways:

- Protect connection endpoints from the possibility that someone could plug an unauthorized device into the network.
- Try to run cable away from general-use and high-traffic areas.
- Secure cable behind walls, above ceilings, and under the floor.
- If a cable run is rarely used, disable switch ports that may connect to it.
- Remove any unused cable runs.
- Label all cables on both ends with the room number, drop number, or other unique ID. Also consider labeling which end device the cable should be plugged into.
- Have a formal diagram or other documentation (switch port, end device, etc.) for all cabling runs when possible.
- Add additional protection to cabling runs that go through physically vulnerable areas, such as break rooms, reception areas, and other high-traffic areas.
- Periodically inspect cabling for any evidence of misuse or tampering.

Wireless media also should be protected to the maximum extent possible. Although there is no physical cabling to protect, wireless access points must be physically protected, and transmission ranges of wireless devices should be limited and controlled. In general, for wireless media, you should

- Use strong authentication and encryption mechanisms (e.g., WPA 2/WPA 3, strong keys, 802.1X, mutual authentication methods, etc.).
- Implement physical protections for wireless access points.
- Use only the power levels needed to transmit to the range of authorized devices; excessive power can send wireless signals into adjoining buildings, the parking lot, and other adjacent areas.
- Centrally locate wireless access points in the facility; try not to place them near windows, external walls, or roofs.
- Filter wireless client MAC addresses when practical, although this is of limited security effectiveness.
- Rename the default service set identifier (SSID) and use SSID hiding whenever possible, although understand that this is also of limited security effectiveness and is more to deter casual wireless snoopers than determined attackers.
- Routinely monitor wireless access points to determine if there is any unusual activity with them, or if there are any unauthorized wireless access points present.
- Segregate guest wireless networks and business wireless networks.

Endpoint Security

Endpoint security is an integral part of the principle of defense-in-depth. Endpoint security operates under the premise that regardless of other security controls and network protections, the host itself must be protected. Whether it is a computer workstation, a laptop, a tablet, a server, or even a smartphone, you should ensure that host devices have thorough security protections, which include the following:

- Strong authentication mechanisms
- Encryption of both data at rest and data in transit
- Limited physical access to only authorized users
- Host-based firewalls and intrusion detection systems
- Updated anti-malware software
- Requirement to connect to corporate networks using secure VPN and NAC
- Device auditing, which includes device access, resource access, and potentially unusual events
- Up-to-date security patches
- Secure configuration settings
- Only minimal applications and services

REVIEW

Objective 4.2: Secure network components This objective summarized measures you should take when securing network devices and related components. We discussed the importance of network architecture and briefly mentioned several key network devices that require physical protection and restricted logical access. Network devices should be protected by using strong authentication mechanisms and encryption; routinely applying security patches; and reducing services and nonsecure protocols. We took a brief look at the different types of firewalls and their functions and discussed network access control devices. We also briefly covered transmission media, both wired and wireless, and some of the security measures you should take to protect it. Finally, we discussed endpoint security and its importance.

4.2 QUESTIONS

1. You are designing an organization's network that must be completely isolated from any outside networks. Only hosts that are locally attached to the network should be able to access resources on it. Which of the following is the best architectural design for the organization?
   - **A.** Internet
   - **B.** Intranet
   - **C.** Extranet
   - **D.** Demilitarized zone
2. You are designing security for a perimeter network for an organization. You are constructing multiple layers of security devices in a demilitarized zone configuration. You want the first layer of security to block unwanted extraneous traffic at the outer perimeter of the network by filtering only on basic characteristics, such as IP address, port, and protocol. Which of the following is the most efficient, least complex device you should use?
   - **A.** Packet-filtering firewall
   - **B.** Next-generation firewall
   - **C.** Proxy
   - **D.** Circuit-level firewall

4.2 ANSWERS

1. **B** An intranet is a closed internal network that is accessible only by internal hosts.
2. **A** A packet-filtering firewall can help offload the effort required by more complex and expensive devices by eliminating a lot of extraneous network noise that is simple to filter, such as traffic that can easily be blocked by IP address, domain, port, protocol, and service. More complex traffic can then be left to be filtered by other devices, such as next-generation firewalls.



Objective 4.3 Implement secure communication channels according to design

This objective concludes our brief discussion of communications and network security. It addresses the key elements that are implemented to protect communication sessions over network infrastructure: identification, authentication, authorization, encryption, data integrity, and nonrepudiation.

Securing Communications Channels

Communications channels are voice, video, and data networks implemented in various ways. In the context of *secure* communications channels, our focus will be the security practices and controls used to protect voice communications and modern multimedia collaboration that occur over networks and the Internet, particularly via remote access. We'll also address communications that take place over virtualized networks, as well as third-party connectivity into sensitive internal networks.

Voice

Voice is considered a "converged technology," in that it has been integrated with network traffic and carried over standard network infrastructures using TCP/IP. From a historical perspective, it's helpful to know that voice only became part of modern networking in the 1990s; prior to that it used the *plain old telephone service (POTS)*, which consists of dial-up lines and dedicated circuits routed through a completely separate system from IP networks. The system was called the *public switched telephone network (PSTN)* and used circuit-switching technology rather than the packet-switching methods used in modern data networks. Regular phone lines used analog systems, and if an organization had its own internal telephone system, voice communication occurred only over what was called a *Private Branch Exchange (PBX)*. These older PBX systems linked an organization's internal telephone system to the outside world, where it was further integrated into the larger public telephone communications infrastructure.

These legacy systems have their share of vulnerabilities, which include generally weak security. These vulnerabilities include

- Lax administration and often unfettered or unrestricted access by anyone within the company, and sometimes even people outside the company
- Weak authentication mechanisms
- No encryption services
- Vulnerability to a variety of attacks, including war dialing, phone freaking, denial of service, and spoofing

Both routine users and hackers used these weak systems to forward calls both into and outside the organization's phone infrastructure and pretend to be organizational members for the purposes of social engineering or to get free long-distance services. Often these systems were also tied to banks of modems, which were used to dial in to an organization from the outside world; that was often the only way people could get connected to their internal network. This allowed war dialing attacks and intrusion into the network through vulnerable PBX systems.

The only mitigations available for these vulnerabilities were to strictly control remote maintenance or administration for PBX systems, allow only limited accounts, limit call forwarding, and configure the device to refuse to allow internal numbers that came from phones outside the organization.

Table 4.3-1 describes legacy voice communications technologies.



EXAM TIP Although considered legacy, you may see questions relating to older voice technologies on the exam.

With better infrastructure and more efficient, reliable devices, voice slowly moved into the data network realm and became part of traffic carried over standard network infrastructures.

TABLE 4.3-1 Legacy Traditional Voice and Dial-Up Technologies

Technology	Description/Characteristics
Dial-up	Connects to the PSTN using a modem (modulator/demodulator), which converts the computer's digital signals and data to analog so it could be sent over regular phone lines to another computer's modem
Digital subscriber line (DSL)	Uses a specialized modem to transmit both analog voice and digital data signals over the regular phone lines to a centralized DSL access multiplexer (DSLAM); comes in two flavors: *symmetric*, where data is sent and received at the same speed, and *asymmetric* (ADSL), most common in residential areas, where download speed is much faster than upload speed
Integrated Services Digital Network (ISDN)	Uses legacy PSTN with specialized equipment to send digital data over analog lines; splits the connection into different channels using three implementations: Basic Rate Interface (BRI), composed of two B channels at 64 Kbps and one D channel at 16 Kbps, used primarily for homes and small offices; Primary Rate Interface (PRI), composed of 23 B channels and one D channel at 64 Kbps; and Broadband ISDN, used primarily as telecommunications backbone connections
Cable modem	Devices providing high-speed access to the Internet using existing cable company coaxial and fiber connections; still in wide use today and uses the international Data-Over-Cable Service Interface Specifications (DOCSIS) for high-speed data transfer

Voice can be carried over standardized IP traffic (called Voice over IP, or VoIP), as can video traffic. To fulfill the high demand for quality voice services in modern businesses, voice traffic requires high bandwidth, resilient connectionless protocols, and availability on a near-constant basis. However, VoIP also inherited some of the vulnerabilities of standard network traffic, such as the risks of interception, eavesdropping, spoofing (i.e., fake calls), and denial-of-service attacks. Unlike secure networking protocols that may have built-in security mechanisms, voice traffic typically has no built-in security protocols to protect it, so it relies heavily on secure networking protocols and devices, the same as other network traffic, to secure it.



NOTE VoIP is also referred to as *Internet Protocol (IP) telephony,* which also includes other technologies, such as real-time messaging and videoconferencing applications.

A variety of VoIP technologies exist, which include both software and hardware. There are also different standards that have been developed over time and have either competed with each other or eventually been merged into each other. Key voice standards that you should be aware of for the CISSP exam are outlined in Table 4.3-2.

VoIP traffic, as mentioned before, is susceptible to the same vulnerabilities and attacks as other types of data traveling on normal IP networks. When implementing voice services in the network, you must consider important issues such as how the organization will make 911 emergency calls and how backup communications will be handled, since the IP network may be susceptible to degradation, interruption, and failure. Backup communications should include wireless or cellular capabilities, as well as traditional PSTN phone systems.

TABLE 4.3-2 IP Telephony Technologies

IP Telephony Technology	Description/Characteristics
Real-time Transport Protocol (RTP)	Session layer protocol that carries data (audio and video) in media stream format; used for VoIP, video teleconferencing, and other multimedia streaming
Secure Real-time Transport Protocol (SRTP)	Enhanced version of RTP that offers authentication and encryption services
Session Initiation Protocol (SIP)	Used to set up and terminate voice call sessions over IP networks; uses a User Agent Client (UAC, an end-user voice application) and a User Agent Server (UAS, a SIP server)
H.323	Standard that facilitates audio and video calls over packet-based networks



NOTE Both SIP and H.323 are competing standards, and both run on top of RTP to provide call session initiation and setup, reliable packet delivery, and other services.

Multimedia Collaboration

Multimedia collaboration is a catchall phrase that refers to a wide variety of collaborative technologies, including video teleconferencing, multiuser application sharing, project workflow, and so on. Although these technologies have been used for a while, they became critical services for conducting business during the COVID-19 pandemic due to the massive numbers of new remote workers who needed to connect to business networks and work with team members on complex projects online. The heightened demand for multimedia collaboration has continued post-pandemic, as many organizations have embraced working remotely as an acceptable alternative.

Some of these technologies, especially video teleconferencing, were not as mature as they could have been prior to the new paradigm of mass remote working, but they quickly caught up with other cutting-edge technologies during this time. Examples of the technologies that became increasingly critical during this increase of teleworking include the following:

- Remote desktop access
- Remote file sharing and collaboration
- Remote meetings
- Collaborative worksites (e.g., Microsoft Teams)
- E-mail and instant messaging

Most of these technologies were originally designed for limited use, such as the occasional remote or traveling user or occasional virtual team meetings. Until the requirement skyrocketed for teleworkers to connect remotely, much of the necessary infrastructure an organization had in place could not support the increased number of people using these technologies. The necessity to rapidly improve collaboration technologies' scalability and performance required infrastructures to follow suit. Many organizations now require more robust, scalable networking infrastructures to support multimedia collaboration on a larger scale; otherwise, they risk serious performance issues, such as speed, latency, and limited bandwidth. Unlike e-mail and web browsing, voice and video technologies require network traffic prioritization to ensure sufficient bandwidth for consistent, stable, no-jitter connections.

As with any emerging technology, the sudden increase in use of multimedia collaboration technologies revealed traditional vulnerabilities, such as lack of authentication and encryption technologies built into the applications and devices, application vulnerabilities that allow them to be exploited over network connections, and administrative and technical issues that go along with remote data sharing and centralized control. These vulnerabilities often result in hijacking attacks, identity spoofing, data interception, and unauthorized remote control of the system or its resources. These types of attacks frequently occurred during the first wave of the

mass teleworking movement as emerging technologies were put to the test in the early part of the COVID-19 pandemic starting in 2020.

The mitigations for these vulnerabilities are the tried-and-true security controls that apply to networking and applications. These include the use of strong authentication and encryption technologies, secure networking protocols, secure application design and implementation, and restrictive access to privileged functions by ordinary users.

Remote Access

Remote network access has been critical for some members of the workforce, such as traveling salespeople, for decades, but, as previously discussed, it has become critical for a broader swath of the workforce since the COVID-19 pandemic began impacting businesses in 2020. Massive numbers of employees are now working remotely, even as many companies are bringing workers back to the offices. We will discuss several different methods for remote access, including old-fashioned dial-up, as you may still encounter it on the CISSP exam, but we will emphasize the more modern methods, such as virtual private networks (VPNs).

Although dial-up connections are no longer widely used in technologically rich countries, you may still occasionally see them being used in rural areas or underdeveloped countries. Dial-up connections require modems, which are used to directly dial into another computer or even to a bank of modems specifically configured for multiple remote access users. Some of these legacy connections may require some form of authentication and accounting server as well, although most of the older versions of these technologies did not use strong authentication or encryption protocols. The older versions of dial-up connections used the now-deprecated Serial Line Interface Protocol (SLIP) or Point-to-Point Protocol (PPP) connections.

Virtually all modern remote access connections use some form of VPN, which allows a communications session that is "tunneled" and protected through untrusted networks, such as the Internet. The untrusted network provides the transport mechanism, while secure protocols are used to provide for both encryption and authentication services. There are two different types of VPNs:

- **Client-to-site VPN** This is used when a single user must connect to a corporate network through a VPN server (also called a *concentrator*) from a remote location, such as a hotel room or home network.
- **Site-to-site VPN** This is used to connect to multiple corporate LANs separated by a wide area connection. It also uses a VPN concentrator but is configured for multiple users to connect across a single local concentrator through the Internet to a remote concentrator, which is part of the organization's infrastructure. Think of a branch office that must connect all of its employees into the corporate LAN through the Internet.

Remote access uses technologies that include protocols and services focused on encryption, authentication, and the connection itself. Remote authentication protocols and services are summarized in Table 4.3-3.

TABLE 4.3-3 Remote Access Protocols

Authentication Protocol/Service	Characteristics
Password Authentication Protocol (PAP)	Older protocol used with dial-up connections that transmits username and passwords in plaintext; no longer used
Challenge Handshake Authentication Protocol (CHAP)	Sends challenge requests and password hashes instead of actual cleartext password
Microsoft MS-CHAPv2	Microsoft's improved version of CHAP
Kerberos	Widely used modern authentication protocol; used extensively in Microsoft Active Directory and other LDAP-based networks
Extensible Authentication Protocol (EAP)	Secure protocol/framework that allows a variety of authentication methods, such as smart cards, Kerberos, and public key infrastructure (PKI)
802.1X	Port-based network access control (NAC) standard that is widely used as an enterprise authentication method in wireless networks, but can also be used extensively in wired networks; allows for multiple authentication and encryption types and can provide for mutual authentication between both users and devices



EXAM TIP You should be familiar with the characteristics of both the authentication protocols/services and the remote connection/tunneling technologies presented here.

In addition to remote authentication protocols, you should be familiar with remote connection and tunneling technologies and protocols, which are summarized in Table 4.3-4.

Cross-Reference

Some of the aforementioned secure protocols, such as Kerberos, EAP, SSH, and SSL/TLS, were covered in more detail in Objective 4.1.

Data Communications

Although an entire book could be devoted to nothing but securing data communications, for the CISSP exam (and real life), you need to be familiar with several key elements: identification, authentication, encryption, authorization, data integrity, and nonrepudiation. To review, this is what these elements mean in the context of data communications:

- **Identification** Communicating hosts and/or users must positively identify themselves before, and even sometimes during, a communications session using identifiers like username/password or multifactor authentication methods (such as PKI).

TABLE 4.3-4 Remote Connection/Tunneling Technologies and Protocols

Remote Connection/ Tunneling Technology	Characteristics
Point-to-Point Protocol (PPP)	Used widely with older dial-up; allows use of plaintext username and password only; works with TCP/IP
Point-to-Point Tunneling Protocol (PPTP)	Advanced version of PPP that uses Microsoft Point-to-Point Encryption (MPPE)
Layer 2 Tunneling Protocol (L2TP)	Modern combination of Cisco Layer 2 Forwarding (L2F) and Microsoft PPTP; no built-in authentication or encryption, only provides tunneling; widely used in today's VPNs
Remote Authentication Dial-In User Service (RADIUS)	Legacy dial-up connection technology primarily used with large ISPs; requires a network access server; supports older protocols including PAP, CHAP, and MS-CHAP; encrypts passwords but no other information; uses UDP ports 1812/1813 or 1645/1646
Terminal Access Controller Access Control System (TACACS, TACACS+)	TACACS, XTACACS, and TACACS+ are three different protocols; they're not compatible with each other. In fact, TACACS+ and XTACACS started out as Cisco proprietary protocols, but TACACS+ has since become an open standard. Uses TCP port 49 and provides authentication, authorization, and accounting functions; uses a variety of encryption and authentication protocols
Secure Shell (SSH)	Secure connection protocol family, uses TCP port 22; provides for mutual authentication and encryption using digital certificates and keys; best used for local/limited remote connections
Secure Sockets Layer/ Transport Layer Security (SSL/TLS)	Secure protocols of choice for HTTPS and WWW implementations; SSL has been deprecated in favor of newer versions of TLS; uses TCP port 443

- **Authentication** Hosts, users, and even network applications or services should authenticate themselves to other entities in the communications process.
- **Authorization** Only properly authorized entities should be able to communicate over the transmission media or during a specific session, to include transmission and reception of sensitive data.
- **Encryption** Both authentication information and data should be encrypted during transmission. Note that encryption can be implemented by encrypting the entire data stream, only certain data within the session, or only the session itself.
- **Data integrity** Technical controls should be in place during communications sessions to ensure data integrity, such as cyclic redundancy checking (CRC), hashing, and other means.
- **Nonrepudiation** Transmitting hosts, users, applications, and services must not be able to deny that they sent a message to other entities in the communications process.

To implement all of these elements to protect data communications, strong security protocols must be used. Many of these protocols were discussed in Objective 4.1. Examples of protocols that should be avoided include those with weak or no authentication or encryption capabilities, such as Telnet, SMTP, FTP, TFTP, and the now-deprecated SSL.

Virtualized Networks

Just as we can virtualize operating systems, software runtime environments, and microservices, we can also virtualize networks. Network virtualization has several advantages, including reducing infrastructure complexity and costs, minimizing the amount of physical hardware, and, of course, in security. From a security perspective, it plays a significant role in sensitive network segmentation and host separation. Networks can use virtual LAN segments, separate IP address spaces, and have different security policies applied to them based on the protection needs of the network. Network virtualization can also be used to dynamically reconfigure a network in case of an attack or disaster; this is something that would be much more time-consuming and difficult using physical hardware devices. We see virtualized networks implemented in several different ways, which are summarized in Table 4.3-5.

Third-Party Connectivity

Ensuring that authorized third parties connect to organizational networks securely is the final category of secure communication channels that you need to understand for the CISSP exam. We have discussed the use of extranets (which can separate sensitive internal networks from those required for external stakeholder access), as well as the availability of network access

TABLE 4.3-5 Virtualized Network Components

Virtualized Network Implementation	Characteristics
Virtual LAN (VLAN)	Creates virtual network segments and subnets using IP addressing; implemented on layer 3 switches; eliminates both collision and broadcast domains; used to separate sensitive network segments; routable using routers and layer 3 switches
Software-defined networking (SDN)	Creates entire logical network architecture using software; separates hardware and configuration details from network services and data transmission; software creates data routing and control infrastructure
Software-defined wide area networking (SDWAN)	Logical extension of software-defined networking; creates software-defined wide area networks
Virtual eXtensible local area network (VXLAN)	Creates virtualized LANs using cloud-based infrastructure; addresses issues with spanning tree, large MAC address tables, and limited numbers of VLANs; tunnels Ethernet traffic over an IP network
Virtual storage area network (VSAN)	Creates a shared storage system over a software-defined network

control (NAC), VPNs, and other remote access technologies to secure connectivity into a network. The same technologies used for an organization's personnel to connect remotely to its network can also allow authorized third parties to connect to the organization's network. These third parties may include business partners, suppliers, service providers, and even customers.

In all cases, you should ensure that the key elements of identification, strong authentication, encryption, authorization, data integrity, and nonrepudiation are present so that only vetted, authorized third-party users can connect to the network. You should also ensure that the same secure design principles used to create secure network infrastructures are used to manage third-party connectivity into the network. These include the concepts of least privilege, defense in depth, secure defaults, separation of duties, zero trust, and so on.

Cross-Reference

The secure design principles were covered in depth in Objective 3.1.

REVIEW

Objective 4.3: Implement secure communication channels according to design In this objective we discussed securing different types of communication channels, including those used for voice, multimedia collaboration, and remote access. We emphasized throughout the objective the use of key security elements, such as identification, strong authentication, authorization, encryption, data integrity, and nonrepudiation. We also discussed how those elements apply to all data communications channels. We then covered the use of virtualized networks, which can separate segments that contain sensitive hosts as well as be dynamically reconfigured in the event of an attack. Finally, we discussed third-party connectivity to internal networks using secure design principles.

4.3 QUESTIONS

1. Which of the following legacy technologies provides symmetric and asymmetric services that affect both upload and download speeds for users?
   - **A.** H.323
   - **B.** Asymmetric digital subscriber line (ADSL)
   - **C.** Plain old telephone service (POTS)
   - **D.** Secure Real-time Transport Protocol (SRTP)
2. Which of the following is a remote access authentication protocol that can use a variety of authentication methods, such as smart cards and PKI?
   - **A.** Password Authentication Protocol (PAP)
   - **B.** Extensible Authentication Protocol (EAP)
   - **C.** Challenge Handshake Authentication Protocol (CHAP)
   - **D.** Layer 2 Tunneling Protocol (L2TP)

4.3 ANSWERS

1. **B** Asymmetric digital subscriber line (ADSL) is a legacy technology that transfers digital data over analog voice lines. With ADSL, more bandwidth is allocated for downstream data than upstream data, so the speeds are different.
2. **B** Extensible Authentication Protocol (EAP) is an authentication protocol/framework that can handle multiple types of authentication methods, including smart cards, PKI certificates, and username/passwords. PAP can only use username and passwords, which are transmitted in cleartext. CHAP cannot use multiple authentication methods. L2TP is a tunneling protocol, not an authentication protocol.



Identity and Access Management (IAM)

Domain Objectives

- **5.1** Control physical and logical access to assets.
- **5.2** Manage identification and authentication of people, devices, and services.
- **5.3** Federated identity with a third-party service.
- **5.4** Implement and manage authorization mechanisms.
- **5.5** Manage the identity and access provisioning lifecycle.
- **5.6** Implement authentication systems.

Domain 5 focuses on the details of identity and access management (IAM). This domain goes into depth on the concepts of identification, authentication, and authorization. Physical and logical access controls are discussed first, followed by the various identification and authentication services used to validate that an entity is actually who the entity claims to be. We will also discuss federated identities with third-party providers, and how authorization mechanisms are implemented after the authentication process is complete. We'll also cover the entire identity and access provisioning life cycle.



Objective 5.1 Control physical and logical access to assets

In this objective we discuss access controls and how they are used to restrict and control physical and logical access to assets, which include information, systems, devices, facilities, and applications.

Controlling Logical and Physical Access

Assets are protected logically and physically using access controls. An *access control* is any security measure that restricts or controls access to resources. *Access* is an interaction between *entities*, which include people, systems, applications, physical locations, devices, and information. The *subject* is the active entity that is attempting to access a resource. The resource is called an *object* and is a passive entity that provides information and services to subjects. Examples of subjects include people, applications, systems, and processes. Examples of objects are applications, systems, and information.

Access controls are designed to identify and authenticate subjects attempting to access resources and determine whether that access is authorized and, if so, to what degree. Certain access controls also monitor and record the interactions subjects have with objects during that access.



EXAM TIP Understand the three major types of access controls, administrative, technical, and physical, as well as the common control functions, which include deterrent, preventive, detective, corrective, and compensating. The types and functions of controls are explained in more detail in Objective 1.10.

Note that for the purposes of the exam objectives, our focus is on controlling logical and physical access to systems, devices, facilities, applications, and, perhaps most importantly, the information that resides in any of these. To quickly review:

- *Information* (including raw data—unorganized facts without context) can be present on systems, devices, and storage media, inside applications, and in other places where information can be transmitted, stored, or processed.
- *Systems* generate, store, process, and transmit information, and include workstations, servers, network infrastructure devices, and even personal or mobile devices.
- *Devices* are a subset of systems, and could include embedded, IoT, and personal mobile devices. Some may have their own unique security requirements to protect them from theft and compromise, such as screen locks, remote wipe, application and information separation (containerization), and a variety of other physical and logical measures.
- *Applications* are software that runs on systems and devices and assists in generating, storing, processing, transmitting, and receiving information. Applications also have their own unique security requirements and access controls, including secure coding, protected resource usage, and so on.
- *Facilities* include multi-building campuses, data centers, server rooms, and any other place where sensitive assets are stored or used. Facilities primarily use physical controls, but logical controls are also useful to connect physical security management systems to the actual controls that restrict and control access to secure spaces.

Rather than break out physical and logical access controls for each of these categories of assets individually, we're going to discuss logical and physical access controls together because they all apply to each of these assets.

Logical Access

Logical access controls include measures usually implemented as technical controls. These controls are most often what we associate with technologies that protect operating systems, applications, and information. As discussed in Objective 1.10, logical access controls are implemented using software and hardware. Key elements used to control logical access include the following:

- Identification and strong authentication mechanisms used to validate and trust subjects
- Authorization mechanisms (rights, permissions, privileges)
- Accountability controls (e.g., auditing)
- Integrity mechanisms (e.g., hashing algorithms)
- Strong encryption systems, algorithms, and keys
- Nonrepudiation mechanisms (e.g., digital certificates)

Physical Access

Many of the controls used to physically protect systems, devices, and facilities are discussed in Objectives 3.8 and 3.9, but they are reiterated here as a reminder:

- Entry point controls to facilities, including sensitive internal areas, such as data centers and server rooms
- Protected "zones" layered within the facility
- Physical intrusion detection alarms
- Video surveillance cameras
- Physical and electronic locks
- Human guards
- Gates, fences, bollards, and other external physical obstacles
- HVAC, temperature, and humidity controls
- Locked equipment cabinets
- Cables and locks for portable systems
- Inventory tags

REVIEW

Objective 5.1: Control physical and logical access to assets In this first objective of Domain 5 we discussed controlling physical and logical access to assets. Assets include information, systems, devices, facilities, and applications. Access is controlled through logical access controls such as strong authentication, encryption, and authorization access controls, and through physical access controls such as locked doors, gates, guards, physical intrusion detection alarms, and video cameras.

5.1 QUESTIONS

1. You must implement logical access controls to protect the operating systems residing on computing devices. First, you need to limit access only to individuals who have been approved to access these devices. Which of the following logical access controls is the first line of defense for controlling access to systems or devices?
   - **A.** Identification and authentication
   - **B.** Authentication and accountability
   - **C.** Identification and authorization
   - **D.** Authentication and nonrepudiation

2. You need to ensure that only authorized personnel access a very sensitive data processing room within your facility. Other personnel can access common areas and other, less sensitive processing areas, so you must set up areas that are progressively more restrictive in nature. Which of the following physical access controls do you need to implement in this scenario?
   - **A.** Physical intrusion detection alarms
   - **B.** Layers of protective zones with differing levels of access
   - **C.** Locked equipment cabinets
   - **D.** Accountability mechanisms, such as system auditing

5.1 ANSWERS

1. **A** The first step in limiting access to a system is to ensure proper identification and authentication of the entity attempting to access it. Unless an individual has been approved for access, they cannot identify themselves and authenticate to a system. Once they are identified and authenticated, then they are authorized, audited, and held accountable for the actions they take, and nonrepudiation is enforced.
2. **B** Setting up layers of protective zones that delineate differing levels of access is a good way to ensure that only personnel approved to enter certain sensitive areas may do so, while also allowing personnel approved for less sensitive areas to access those areas.



Objective 5.2

Manage identification and authentication of people, devices, and services

In this objective we examine key concepts related to managing the identification and authentication of entities or subjects, which can include people, devices, applications, services, and processes.

Identification and Authentication

Access control is composed of several important components, including identification, authentication, authorization, and accountability. We will discuss identification, authentication, and accountability in this objective (as well in Objectives 5.3, 5.5, and 5.6), and discuss authorization in depth in Objective 5.4.

In the context of access control, *identification* is the process of presenting proof of identity and asserting that it belongs to the specific entity presenting it, such as a user. Identification alone does not grant an individual entity anything. The identity must be validated (proven to be tied to the entity presenting it) by some mechanism that is trusted. This is the essence of authentication; an entity asserts a specific identity, and that identity is validated as true through an authentication process. Authentication works by examining the credentials the entity provides, which often consist of a username and password, or other identifying information, and verifying in a centralized system, usually a database of some sort, that what the user has supplied to assert that identity is contained in the user database and is trusted. It requires the user to assert information that could only be found in the trusted database.

Identity Management Implementation

Identity management (IdM) is a formalized, ongoing process that an organization conducts for all entities accessing its information systems. IdM is crucial for anyone who manages resources, such as those on the Internet, that are accessed by customers, suppliers, business partners, and other users. An organization may have their own internal IdM solutions in place, but many organizations outsource this function to third-party providers. This is where third-party or federated identity management comes into play, which we will discuss later in this objective.

IdM requires managing identities throughout their entire life cycle. This life cycle includes identity creation; secure key and identity storage; and use—updating, suspending, and finally, when their time has come, revoking identities. This is closely tied to the identity provisioning life cycle, discussed in Objective 5.5, but IdM really consists of everything we will discuss in this domain.

Single/Multifactor Authentication

In order to authenticate, a user must provide credentials. Credentials could consist of a username and password, smart card, token, or even a fingerprint and a personal identification number (PIN). These credentials, or identifying characteristics, can be categorized as *factors*.

There are different factors you can use for authentication. These include, among others:

- Something you know (knowledge factor)
- Something you have (possession factor)
- Something you are (inherence factor)

These factors are used singularly or in different combinations to provide identity to authentication mechanisms. If you only use a single one of these factors, such as the aforementioned username and password combination, this is called *single-factor authentication*. For example, a username and password are both something you know, and can be provided as a credential. Although it's two separate pieces of information, it still considered one factor, in this case, the knowledge factor.

Requiring more than one factor is called *multifactor authentication (MFA)*. An example of multifactor authentication is the use of an ATM card, smart card, or token combined with a PIN. In this case, two factors are involved: something you possess (the card or token) and something you know (the PIN). Inherence factors would include fingerprints, voice prints, facial recognition, and handwriting, since these are factors that are inherent and unique to an individual. Note that inherence factors are also called *biometric* factors, since these help make up who you are and are uniquely identifiable.



NOTE In addition to the three main authentication factors discussed in this objective, other factors that can be used in MFA include location and time of day and other rule-based characteristics that make multifactor authentication even more secure.

Multifactor authentication (also referred to as *strong authentication*) is defined as using multiple factors for identification and authentication, because compromising two or more factors is much harder than compromising a single factor. For instance, it would be relatively easy to compromise a username and password combination (single factor) since an attacker could discover both of those. However, in the case of a smart card and PIN combination used for multifactor authentication, even if an attacker had one of them, they might not have access to the other.

An organization should implement MFA for all systems that hold valuable information or are connected to the organization's network. MFA is becoming the preferred method of authentication in businesses, on websites, and even for personal use, but there is still widespread use of single-factor authentication, primarily the simple username and password combination. Even with password complexity requirements, passwords can still be compromised with relative ease.



EXAM TIP Remember that multifactor authentication requires at least two of the aforementioned types of factors, such as a fingerprint (inherence) and PIN (knowledge). While a username and password combination includes two distinct pieces of information, it consists of only the knowledge factor and is not considered multifactor authentication.

Accountability

Accountability is one of the tenets of security we discussed early in Domain 1 (in Objective 1.2). Accountability ensures that the actions an entity performs can be traced back to that entity, and the entity can be held accountable for those actions. Accountability is closely related to nonrepudiation, which ensures that an individual cannot deny that they took an action. Accountability is made possible by auditing, another closely related security tenet, as auditing records the interactions a user has with systems and resources.

Accountability is a key component of identity and access management and is built into many of the technologies we will discuss, including authentication and authorization protocols and technologies. The key element to understand about accountability is that users' and other entities' interactions with systems and resources are recorded, including all details about the authentication process they undergo as well as their use of authorized privileges.

Session Management

All communications sessions between two parties, such as a user and a website, or a user and the network during the day at work, must be managed and secured for the duration of the session. The session starts with identification and authentication, and usually ends with termination of the session with an application, a system, or even a network. During a session, data is exchanged that must be protected, including credentials used for periodic reauthentication. Additionally, the session itself may be subject to security requirements, such as a limited length of time, conducted between only specific hosts, and so on. The following are some key security measures you should ensure are implemented to provide for secure session management:

- Strong authentication mechanisms
- Strong session keys
- Strong encryption algorithms and mechanisms
- Timeouts to limit session duration
- Session inactivity disconnection
- Controls to detect anomalous activity (e.g., activities associated with session hijacking)

Registration, Proofing, and Establishment of Identity

Authentication means that an individual's system credentials are validated as being tied to that individual, so the system can verify that they are who they say they are. However, an individual cannot be authenticated by a system unless that system already has the individual's verified credentials. But how is the individual validated before they even get an account? It requires a process whereby the individual *establishes* their identity with another entity and validates that identity. *Registration* and *identity proofing* are interrelated processes organizations use to prove the identity of an individual, as well as to register that identity for future validation. Once that identity is validated, a chain of trust between it and other identities, such as system credentials, can be established.

Take, for instance, the process of getting a U.S. passport. A passport itself is a valid proof of identity, but how do you establish your identity in order to get a passport? You have to provide other trusted forms of identification, such as a driver's license and birth certificate. You also have to appear in person before a passport agent and provide an approved picture and likely have your fingerprints made. This establishes a chain of trust so that the passport can be issued, which itself is trusted afterwards. This process is called *establishment of identity*.

For new employees in a company, establishment of identity involves providing physical documentation, such as a driver's license, birth certificate, passport, and so on, to the employer. Employers require this information for various reasons, including to initiate security clearances, withhold taxes, and perform background checks. This process must be trustworthy; if an organization mistakenly accepts proof of identity that is not actually valid (such as a library card or an expired driver's license) or identification that could be easily forged, that sets up a chain of weak trust and proofing throughout the onboarding process as well as other processes.

This proofing and establishment of identity process is used not only by employers, but also by any organization that owns a resource a user might wish to access online. Think of online banking, for example, or access to another financial website. The individual must prove their identity, often by completing online forms that verify who they are through a series of questions with answers that only they would know, or by providing, again, written documentation in person before being allowed access to the site. In any event, the registration, identity proofing, and establishment of identity process is used to validate the individual so that the organization can subsequently issue the individual a trusted set of credentials.

Federated Identity Management

Federated identity management (FIM) means that multiple organizations may use a single IdM provider. The IdM provider stores individual identities and credentials belonging to an entity which are validated upon request to any other entity wishing to authenticate an individual, across multiple organizations participating and systems. For example, Facebook performs FIM services when users wish to use their Facebook credentials to authenticate to a third-party application or website. More often than not, a third party is responsible for federated identity management, but the key here is that the same identity, once established, can be used for various resources, regardless of the organization, as long as the organization trusts the federated identity provider's credentials.

Cross-Reference

Objective 5.3 continues the discussion of federated identity management with third-party services.

Credential Management Systems

A credential management system is one that securely stores credentials and can be used to issue, update, suspend, and revoke credentials. Consider an Active Directory account that can be created, added to groups, granted permissions, locked, and deleted. The Active Directory database, along with the utilities used to manage the user accounts in the database, is a form of a credential management system. However, credential management systems are much more comprehensive, as they can also include the administrative policies used to manage credentials, the physical systems that control and use smart cards, and other technical controls.

In addition to creating, updating, and deleting accounts, credential management systems facilitate the management processes used in long-term account management. Some of these management processes are embedded in enterprise credential management systems, but others may be add-ons or even standalone applications. Other activities performed by credential management systems include

- **Password management** Credential management systems assist in the creation and secure storage of user passwords
- **Password synchronization** Credential management systems assist in synchronizing passwords across multiple systems, applications, and identities.
- **Self-service password reset** Credential management systems assist users in resetting their own passwords by allowing them to use cognitive passwords (e.g., a series of predetermined questions and answers that only the user would know), or other secure methods.
- **Assisted password reset** Credential management utilities can be used by helpdesk personnel to assist in resetting user passwords securely, allowing the user to validate their identity and create a password that only they know.

Single Sign-On

Single sign-on (SSO) is a method of authentication that requires a user to only authenticate once and still have access to multiple resources within an organization and possibly even within other organizations. Before single sign-on, users had to authenticate several times if each resource had its own security policies and credential requirements. Single sign-on is implemented by various technologies, including Kerberos and Windows Active Directory. Other technologies can assist in integrating single sign-on with federated identity management so that third-party identity providers can be used to authenticate a user to resources spanning multiple organizations, websites, and so on.

Cross-Reference

Objective 5.6 discusses technologies that can assist with single sign-on authentication across multiple organizations, including OpenID Connect (OIDC) and Open Authorization (OAuth), as well as provides an in-depth discussion on Kerberos.

Just-in-Time

Just-in-time access control refers to only allowing a user to have the access they need at the time they need it. Usually JIT access applies only to a specific set of actions and is temporary. For example, running a privileged command as a nonprivileged user may require an employee to escalate their privilege level by authenticating as a different account with higher privileges. In Windows this is often accomplished by using the `runas` command, and in Linux by using

the `sudo` command. Just-in-time access is usually provisioned ahead of time but does not become effective until the user actually needs it. Access may also be contingent on circumstances; a user may not have access to a particular resource or a specific privilege level unless they need it based upon circumstances, such as time of day, login host, and so on.

Cross-Reference

Objective 5.5 provides details about the `runas` and `sudo` commands.

Just-in-time access prevents users from carrying higher privileges or access than they normally need on a continual basis; they only get the access when they need it and only under specific circumstances. Such access can also be temporary and apply to only a specific action or resource. Just-in-time access is an effective way to adhere to the principle of least privilege while still allowing users to have the functionality they may need based upon their job duties. JIT access also helps to relieve the IT administrative burden by enabling users to perform necessary tasks without IT intervention.

REVIEW

Objective 5.2: Manage identification and authentication of people, devices, and services In this objective we addressed identification and authentication of people, devices, and services. We discussed several key elements of identification and authentication, including identity management (IdM), authentication factors (knowledge, possession, and inherence), and the distinction between single-factor authentication and multifactor authentication.

We also discussed the accountability aspect, which means that details about the identification and authentication process must be recorded in order to hold individuals accountable for their actions. Session management is ensured through strong identification and authentication mechanisms, as well as strong encryption.

In order to initially prove an individual's identity, they must go through a registration, proofing, and establishment process. Once this process is complete, the individual is able to obtain other identities, such as system credentials, that can be further trusted.

Federated identity management uses a single entity that is responsible for credentials that multiple organizations can trust and use to authenticate an individual for multiple resources. Credential management systems are used to create, update, and revoke credentials, as well as perform functions such as self-password reset, password synchronization, and password management. Single sign-on is an implementation of authentication technologies that allows the user to authenticate only once but subsequently be able to access many different resources.

Just-in-time access control refers to only allowing a user to have the access they need at the time they need it. Usually JIT access applies only to a specific set of actions and is temporary. The use of the `runas` command in Windows and `sudo` in Linux can facilitate JIT access control.

5.2 QUESTIONS

1. You are tasked with implementing multifactor authentication for a sensitive system in your organization. Your supervisor would like you to explain which combination of factors would be considered multifactor. Which of the following would you tell your supervisor is a good example of multifactor authentication?
   - **A.** Smart card and PIN
   - **B.** PIN and password
   - **C.** Username and password combination
   - **D.** Smart card and token
2. In order to issue you credentials for a sensitive system, your organization must go through a formal process to verify your identity. Which of the following examples best describes registration, identity proofing, and establishment processes?
   - **A.** Providing your public key to your employer
   - **B.** Proving your identity using a driver's license, passport, and birth certificate
   - **C.** Providing a letter of recommendation from a supervisor
   - **D.** Comparing the picture on your identification card to a picture in a company database

5.2 ANSWERS

1. **A** Multifactor authentication consists of at least two separate types of factors. In this case, a smart card, which is something you possess, and a PIN, which is something you know, would be two factors. All other choices would be considered single-factor authentication, since each one of them only uses one of the authentication factors.
2. **B** Registration, identity proofing, and identity establishment processes require that you conclusively prove who you are, using trusted identification credentials. In this case, a driver's license, passport, and birth certificate would conclusively prove your identity. Providing your public key to your employer would not prove anything, since anyone can have your public key. Providing a letter of recommendation from a supervisor would not prove who you are. Comparing the picture on your identification card to a picture in a company database would not conclusively prove who you are, since both the identification and the information in the database could be forged.



Objective 5.3 Federated identity with a third-party service

In this objective we will examine federated identity management (FIM) processes that use third-party services. Third-party services are most often used in online identification and authentication, and examples include Microsoft and Google authentication services and apps that people have on their smartphone or tablet. This objective explains third-party identity management services in on-premise, cloud, and hybrid implementations.

Third-Party Identity Services

Not every organization can afford the cost of the personnel, infrastructure, and policy involved with identifying and authenticating external individuals who access the organization's web services and managing their credentials. Fortunately, such organizations can turn to a third-party identity service provider to meet their needs, which often is more affordable and more efficient and ensures that identification and authentication are securely handled.

Almost everyone has used third-party identity services, whether they actually realized it or not. If you've ever used Duo, Microsoft, or Google authentication services, for example, you have used third-party identity services. For this objective, we will discuss three approaches to providing third-party identify services: on the premises of the organization, through cloud-based services, or as a hybrid of both methods. In Objective 5.6 we will discuss additional components of identity and authentication services, including mechanisms and protocols associated with third-party services such as the Security Assertion Markup Language (SAML), Open Authentication (OAuth), and OpenID Connect (OIDC).

As briefly discussed in Objective 5.2, FIM expands beyond a single organization and is used when organizations join a group or a "federation" of other organizations who share identity and credentials for the purposes of single sign-on (SSO) into multiple resources across organizations. Examples of a federation might include a company with numerous geographically separated divisions or campuses, universities with multiple campuses, or government organizations across a country. Note that while the federation manages identities and credentials, each member of the federation controls its own access to resources.

On-Premise

Having an on-premise (aka on-premises) solution for identity management (IdM) typically means that the organization has implemented and retains both management control over the solution and the responsibility for its daily maintenance and security. However, some third-party solutions also can reside on-premise; these are solutions that require identity management servers or services to be installed within the organization's internal infrastructure.

The organization still has some level of control, as well as lower-level maintenance and security responsibilities with the solutions. These are typically integrated solutions that use standardized identification and authentication mechanisms and protocols, such as Kerberos authentication and Microsoft Active Directory.

Cloud

Cloud-based identity services are third-party Identity as a Service (IDaaS) offerings that provide identification, authorization, and access management. This scenario is most often encountered when an organization's internal clients use Software as a Service (SaaS) applications from a cloud service provider. A common example is the use of Microsoft Office 365 and Azure cloud-based subscription services, which can be integrated with Active Directory. Note that cloud-based solutions may require a cloud access security broker (CASB) solution, which controls and filters access to cloud services on behalf of the organization. A CASB can provide access control, auditing, and accountability services to a variety of cloud-based services. Additionally, cloud-based services can allow for resiliency and higher availability services, since there are multiple redundancies built into the cloud provider's data center. Cloud-based services, compared to on-premise ones, also offer significant cost savings since the organization is not required to acquire or maintain its own equipment, nor retain the trained personnel needed to maintain it.



EXAM TIP Understand the advantages and disadvantages of both on-premise and cloud identity services. On-premise services offer more control for the organization, but cloud services offer resiliency and cost savings.

Hybrid

Hybrid identity services solutions integrate both cloud-based and on-premise IdM solutions. A hybrid model gives an organization the best of both worlds: it allows the organization to retain a certain level of control while at the same time offering the benefits of cloud-based solutions, such as resiliency and cost-effectiveness.

REVIEW

Objective 5.3: Federated identity with a third-party service In this objective we examined three approaches to implementing federated identity with a third-party service. With an on-premise solution, the organization integrates a third-party identity provider's services into the organization's on-premise infrastructure. Cloud-based solutions are offered as a service from cloud service providers. Hybrid solutions are a mixture of both on-premise and cloud-based solutions and offer the best of both worlds in terms of organizational control, resiliency, and cost-effectiveness.

5.3 QUESTIONS

1. If an organization wishes to maintain strict control over its identity and authentication mechanisms, using its own infrastructure and resources, which of the following is the better solution to ensure that it maintains control and security of those services?
   - **A.** Hybrid
   - **B.** Cloud
   - **C.** On-premise
   - **D.** Federated identity management (FIM)
2. Your organization requires identification and authentication services separate from its on-premise infrastructure, but only specifically for certain applications provided through a Software as a Service (SaaS) subscription. You have enabled pass-through authentication from your on-premise solution so that those credentials can be passed through a CASB for single sign-on capabilities for your software subscription. Which of the following types of IdM solutions have you implemented?
   - **A.** On-premise
   - **B.** Cloud-based
   - **C.** Federated
   - **D.** Hybrid

5.3 ANSWERS

1. **C** On-premise solutions are the best for organizations that wish to retain strict control over their IdM solutions. This grants them exclusive control over their solutions, but also is more costly in terms of the infrastructure they must implement and the personnel required to maintain the systems.
2. **D** Since it contains elements of both on-premise and cloud-based solutions, this is a hybrid setup.



Objective 5.4 Implement and manage authorization mechanisms

This objective builds upon our discussion of access control in Objective 5.1 by examining authorization. After a subject has been identified and authenticated, authorization determines which objects the subject is permitted to interact with. In this objective we will discuss access control concepts, as well as revisit access control models.

Authorization Mechanisms and Models

Validated identification and authentication are the first steps in allowing access to objects by subjects. However, just because a subject has been identified and authenticated does not mean the subject has carte blanche access to everything. Authorization mechanisms ensure that authenticated entities can only perform the actions that they are allowed to perform, and no more. There are several mechanisms and access control models used to ensure this restriction. Note that authorization models and mechanisms are designed around the core security principles of least privilege (an entity should only have the minimum level of rights, privileges, and permissions required to perform their job duties, and no more) and separation of duties (no individual should have more access than they need such that they are able to perform critical tasks alone).

Cross-Reference

The security principles of least privilege and separation of duties are covered in depth in Objectives 3.1 and 7.4.

The following are key concepts related to access control that you need to be aware of for the CISSP exam:

- *Security clearances* are levels of trust based on background checks and other factors that verify an individual can be trusted with a particular level of sensitive information.
- *Need-to-know* means that an individual requires access to the system or information in order to perform the tasks required by their job role.
- A *constrained interface* assists in restricting the actions a subject can take with an object by controlling which actions they can perform through the operating system or application.
- *Content-dependent access* means that users are restricted based on the type of information that an object holds, such as healthcare or financial information.
- *Context-dependent access* means that users must be in the correct environment and context in order to perform specific actions; for example, a user may only be able to access administrative functions from a specified host, such as a jump box, and no other host.
- *Permissions* are the allowable actions that can be taken on a specific object, such as reading or writing to a file or shared folder. Note that permissions are characteristics of the object, not the subject.
- *Rights* refer to the capabilities that entities have, such as to restore backup data, that routine users do not have and must be explicitly granted.
- *Privileges* is often used as a catchall term that refers to special actions that a routine user cannot take, with regard to both systems and data.



NOTE The terms permissions, rights, and privileges are often used interchangeably but have subtle differences. Permissions normally refer to actions that are inherent to an object that can be granted to users. Rights refer to capabilities that must be specifically granted to a user but are not necessarily object dependent. Privileges can mean either of those two things, depending on the context and the specifics of the interaction between subjects and objects, but always refer to actions or capabilities that must be specifically granted to an entity and are not normally part of a routine user's abilities.

The granular requirements that we will discuss for granting or denying access by subjects to objects can be either explicit or implicit. *Explicit* means a permission, approved action, or a rule is specifically identified or listed as allowable. *Implicit* means that something that is *not* specifically listed may still may be implied or effective by default. For example, a user could be explicitly granted write permission to a folder, and that permission is identified in the folder's access control list. However, if the user is granted full control of a folder by the folder's owner, then that user implicitly has the permission to write to that folder, even though that specific permission was not explicitly listed. Explicit and implicit permissions and rules can be tricky to navigate, especially when it comes to rule sets in firewalls, routers, and other security devices.

In the remainder of this objective, we will examine six authorization models: discretionary, mandatory, role-based, rule-based, attribute-based, and risk-based access control.

Discretionary Access Control

Discretionary access control (DAC) means that the creator or owner of an object (a resource such as a file or folder, for instance) can assign permissions and other access controls to the objects they manage. Even a routine user can be the creator or owner of an object and assign permissions to others. Depending on how the DAC system is set up, such as in Windows, administrators can also assign those access controls in addition to creator/owners. This is the most common model you will see implemented in modern operating systems.



EXAM TIP DAC is the only form of access control model that is discretionary in nature; in other words, object creators and owners can grant or deny access to the object in question. The other models we will discuss are nondiscretionary and access control decisions are managed by security administrators.

Mandatory Access Control

Mandatory access control (MAC) means that the security administrator must assign all access, and creators or owners of objects do not have the ability to assign permissions, rights, or privileges at their whim. Mandatory access control is often implemented in highly secure systems,

such as those used by the government or medical fields, so that only authorized subjects can access very specific objects. Mandatory access control specifies requirements such as:

- Security clearance level held by individual
- Need-to-know required by job position
- Explicit supervisor approval for access
- Security labels for objects that match security clearance

Cross-Reference

Mandatory access control was also discussed in Objective 3.2 and is used in confidentiality and integrity models such as Bell-LaPadula, Biba, and Clark-Wilson.

Role-Based Access Control

Role-based access control (RBAC) is also nondiscretionary in nature—the security administrator assigns all permissions, rights, and privileges based on the user's role or job function. For example, a database operator role may only have specific permissions to manipulate objects and data within a database, and nothing else. When a user is placed in that role, they inherit those specific permissions. If the user is filling multiple roles, such as both database operator and supervisor, they may accumulate permissions unless this is specifically denied by administrators.

Note that roles and groups are not necessarily the same thing; those terms are often used interchangeably, but groups are logical collections of users, computers, or other objects that can be assigned or inherit certain elements of security policies. Roles are defined collections of users with specific required privileges based on their job function. Although modern desktop operating systems and networks have built-in roles and groups, these are not the same thing as true RBAC.



EXAM TIP Although there are role-based, rule-based, and risk-based access control models, it is accepted convention in the security world that the RBAC acronym is associated with role-based access control only. If you see the term RBAC on the exam, it most likely refers to role-based access control.

Rule-Based Access Control

Rule-based access control models are used in specific instances; the most common examples are firewall and router rule sets. Rules are made up of specific elements and conditions, and the actions that match those rules are either denied or allowed. Note that rules can be applied to other models as well; DAC can have rules imposed on users, such as specific login hours or IP subnetworks.



EXAM TIP Rule-based access control is almost always found in traffic-filtering devices, such as firewalls and routers; however, rule-based access control can also be integrated into other models.

Attribute-Based Access Control

Attribute-based access control (ABAC) is very similar to rule-based access control but goes into a much deeper level of detail. Attributes used for access control may include almost any detail regarding users, hosts, networks, and devices. Access can be controlled by any number or combination of precise attributes, such as day of the week and time of day, login location, authentication mechanism used, and so on. Access control attributes can be based on the subject or the object, or a combination of both. For instance, attributes like the user's login host and the characteristics of a shared resource, such as sensitivity level, can be used in combination to allow or deny access.

Risk-Based Access Control

Risk-based access control is dynamic in nature; this is also a newer model that does not use fixed requirements or characteristics, as does traditional MAC, DAC, and RBAC. This model adapts to the current situation and changing characteristics of system and environmental risk. The model learns and applies a previously defined baseline of subject and object behavior, and then examines subsequent behaviors and the current risk environment to determine if there is an increased risk for allowing a subject to take an action. Note that risk-based access control still uses defined policies and rules but can dynamically change to allow or block access based on a changing risk environment.



NOTE Although we describe distinct access control models here, in reality they are often combined and used together in modern access control systems. You will likely find DAC, RBAC, and rule-based access control used together frequently, albeit on different systems or applications in the infrastructure. You may also find instances of MAC used with role-based and rule-based access controls.

REVIEW

Objective 5.4: Implement and manage authorization mechanisms In this objective we discussed authorization and access control models and mechanisms. We reviewed key concepts of access control, including need-to-know, validated identification and authentication, and security clearance. We looked at access control mechanisms such as constrained interfaces, content-dependent access, and context-dependent access.

We also defined permissions, rights, and privileges. We then examined the six different authorization models, including discretionary access control, mandatory access control, role-based access control, rule-based access control, attribute-based access control, and risk-based access control.

5.4 QUESTIONS

1. You are a security administrator in an environment that processes sensitive financial information. One of your users calls and asks if you can grant read access to a shared folder containing nonsensitive information to another user in her department. She is the folder's owner, so you explain to her how she can give permissions for access to the folder herself. Which of the access control models is in use in this environment?
   - **A.** Nondiscretionary
   - **B.** Discretionary
   - **C.** Mandatory
   - **D.** Role-based
2. You are a security administrator for a large financial company and you are setting up access control for a new system. You must retain the ability to control which personnel get access to the system, and individual accounts are not permitted to have access. You are going to create defined categories of users, each with very specific access to the system based on their job function. Which of the following access control models best meets your requirements?
   - **A.** Rule-based access
   - **B.** Discretionary access
   - **C.** Risk-based access
   - **D.** Role-based access

5.4 ANSWERS

1. **B** Since the owner of a shared folder can grant access without requiring intervention of a security administrator, this is the discretionary access control (DAC) model.
2. **D** The model that best meets your requirements in this scenario is role-based access control, since you are creating categories of users (roles) that will be assigned access based on their specific job role.



Objective 5.5 Manage the identity and access provisioning lifecycle

This objective discusses the identity and access provisioning process. Since this is a defined, repeatable process with standardized steps, we can organize its activities into a life cycle. As with all life cycles, there is a beginning and end, and it doesn't just begin with giving someone a user account or end by simply deactivating that account when they leave the organization.

Identity and Access Provisioning Life Cycle

Provisioning is the set of all activities that are carried out to provide information services to personnel. Whether the information service is providing access to an application, permissions to a shared folder, a new user account, or even a new laptop, there is a defined life cycle for this process. The process covers the gamut of creating accounts, managing those accounts and the resources they have access to, periodically changing the accounts, and finally terminating those accounts when they are no longer needed. The identity and access provisioning life cycle must be based on several key security concepts, which include the principles of least privilege, separation of duties, and accountability. This process lends a framework to the major steps of identification and authentication, authorization, auditing and accountability, and nonrepudiation.

In the following sections we'll discuss the major pieces of the identity and access provisioning life cycle, including account access provisioning and periodic review, formal provisioning and deprovisioning processes, placing people in the correct roles based on their job functions, and handling on-the-fly privilege escalation.

Provisioning and Deprovisioning

Provisioning and deprovisioning information resources for users consists of many different steps that occur at various stages, including when an individual is hired, during the term of their employment, and upon employment termination or retirement. Many of the provisioning steps are repetitive and many of the deprovisioning steps are simply the reverse of the provisioning steps that created the account or granted the access. We will look at three phases of this life cycle: first at initial provisioning, then at the deprovisioning process that occurs when users are terminated or retire. Later in the objective we will examine the activities, performed in between these two events, that make up managing access while the user is still employed, such as defining roles, privilege escalation, and periodic account review.

Initial Provisioning

Most organizations have an onboarding process, which typically involves briefing new hires on their first day about benefits, company policies, and security responsibilities. This activity

almost always includes the IT department provisioning accounts for new employees. Simply creating an account from an IT perspective is quite easy, but there is some background work that should be done beforehand.

During the initial onboarding, and likely prior to final employment decisions, individuals usually undergo a background check. This could include lower-level background checks such as verification of identity and citizenship, a simple credit check, a criminal background check, and so on. Some background checks are more intense, such as the ones associated with government security clearances. Often, when individuals transfer from one job role to another within the organization, much of this background check information follows them, as is the case with government security clearances, so the organization does not have to start over from scratch. Note that IT does not have any involvement in this portion of on-boarding; these are strictly human resource department functions. However, they inform the level of access the individual is granted to organizational information assets.

Assuming the background check passes muster and the individual meets the other qualifications, which may include minimum education requirements, minimum experience, particular certifications, and so on, the individual is hired and brought into the organization. Since individuals are hired for particular job functions, the organization should have a process in place to obtain approvals for access to various levels of information sensitivity prior to the individual's first day on the job. This may involve coordination between the functional supervisor or data owner and the human resources and security departments.

Once this coordination is completed, the IT department creates the user account and assigns the appropriate access to various resources, as determined by the individual's job position. The user account creation process is often called *registration* or *enrollment* and normally consists of assigning a unique identifier to the individual and provisioning various authentication methods, such as multifactor authentication tokens, smart cards, and fingerprint enrollment. During this onboarding process, the user is always expected to sign a document indicating that they have been briefed on and understand policies such as acceptable use, nondisclosure, and so on.

Deprovisioning

Users may be deprovisioned for several different reasons. As a simple example, a company that loses a government contract would be required to ensure that its employees turn in their government-issued access badges and that their accounts are removed from the systems that apply to that contract. Other events that prompt the deprovisioning process include off-boarding activities such as retirement and termination, which result in leaving the organization. Note that while IT personnel have limited to no involvement in the personnel activities involved with retirement or termination, they have responsibilities that focus on deprovisioning access to information systems in the organization.

The following are the most common deprovisioning activities that IT personnel are likely to be involved with:

- Removing user accounts from systems, data storage areas (such as shared folders or databases), and applications
- Backing up any critical data that may have been owned by the user account

- Decrypting any encrypted data the user may have
- Requiring the user to turn in equipment, access badges, tokens, and other identifiers
- Suspending the user's accounts



NOTE It's generally not a good practice to delete a user's account immediately after the user leaves the organization. You'll almost always find data that you didn't know existed or that needed to be decrypted, and you may not be able to access that data after the account is deleted. Standard practice after a user has left the organization is to only *suspend* the user's account (by deactivating it or simply changing the password) for a predetermined length of time, such as 30 days, and then delete it.

If a user is terminated, particularly for cause, the organization must be very careful in how it handles deprovisioning. The user may or may not know they are being terminated, and the organization must make sure that they don't have time to perform any actions with information or systems that could be detrimental, such as unauthorized copying or deletion of data, damaging systems, encrypting massive amounts of critical data, introducing malware into systems, and so on. As soon as the decision is made to terminate an individual, their user account must be immediately suspended so they no longer have access to any system, application, or information. They should be escorted at all times by someone in management or security. All of their access tokens, badges, and other identifiers and authenticators must be immediately confiscated, and their equipment must be immediately turned in so that the organization can later retrieve data from the devices.

Retirement or any prolonged absence from the organization (e.g., extended vacations, schooling, and sabbaticals or even longer-term disability) may not require the stringent requirement of immediate account termination and escort throughout the facility. Usually these types of termination are not acrimonious, so once the organization is aware of the individual's pending exit, management can decide whether to allow the individual to work for a specified period of time (such as when the individual gives a 30-day notice) based on the trust the organization has in the individual. Regardless, the same process is followed on the individual's last workday; equipment is collected, access is revoked, badges and access tokens are turned in, and the individual is escorted out of the facility.

Role Definition

Roles should be defined and assigned to users when they are initially provisioned; however, those roles will likely change over the long-term course of their employment. Transfer to another department or project, promotion within the same department or to a different one, or changing job function and responsibilities entirely all may necessitate changing the roles assigned to users.

Remember that roles (and even groups in a discretionary access control model) are created so that individual user accounts do not need to be assigned privileges. The user accounts are simply placed in the appropriate predefined role, such as database operator, supervisor,

administrator, financial controller, and so on. This makes managing access privileges much simpler for administrators than it would be if individual user accounts were assigned permissions, rights, and privileges. Continually updating role membership is a part of the provisioning and deprovisioning life cycle.



EXAM TIP Note that role-based access control (RBAC), discussed in Objective 5.4, does not necessarily have to be implemented in order to assign people to proper roles. Even in discretionary access control models, users can be assigned to appropriate roles or groups, which are assigned the proper rights, permissions, and privileges.

Privilege Escalation

Just as an individual can change roles throughout their career, they can also change privileges. Often changing roles and privileges go hand in hand; a database operator who is promoted to a supervisor within the department may retain the role and privileges of a database operator, for instance, but also accumulate additional privileges that come with the role of a supervisor. An individual can be granted additional privileges by simply adding that individual to another role or group. They can also be granted additional privileges by individual user account, although this isn't the preferred method since the privileges can be more difficult to track, resulting in *privilege creep* (the user accumulates excessive privileges over a long time).



CAUTION Even when considering privilege escalation, you should still make the effort toward minimizing additional privileges whenever possible or practical and adhere to the fundamental security principle of least privilege.

When a user requests additional privileges, a supervisor or someone else with decision-making authority should verify that the escalation of privilege is necessary for the user to perform their job. Privilege escalation should not happen simply as a result of a user contacting the IT support desk and demanding additional privileges. Granting additional privileges should also be documented and carried out using roles or groups whenever possible.

Temporary Privilege Escalation

In some situations, a user may need privilege escalation only temporarily, such as to install a new piece of software or run a system update. This is the basis of just-in-time (JIT) access provisioning, introduced in Objective 5.2. Since users should not have constant privileges that are above what they routinely need to do their job, exceptions could be provisioned by simply granting them the ability to run privileged functions for a short period of time on the fly.

In Windows, temporary privilege escalation is accomplished using the `runas` command, which a user can access by right-clicking the icon of an executable program, provided the user has another privileged account they can access to run the program. This approach serves several purposes. First, it prevents the user from being constantly logged in with excessive privileges that they could misuse. Second, it requires just-in-time authentication (as the user must input the required credentials for the privileged account). Third, the account can be audited more stringently so that these actions are recorded for review.

In Linux, just-in-time privilege escalation can happen using the `sudo` command, provided the user is included in a predefined *sudoers* configuration file. The file also includes an entry stating which privileged group the user can leverage and which commands they can run. The `sudo` command also requires that the user authenticate when they are attempting to use the utility, so that their actions are auditable. The use of the `sudo` command to temporarily escalate privileges removes the need for the user to log in as the system root account.

Managed Service Accounts

Managed service accounts are those that don't belong to normal users; they belong to applications or processes that must execute programs automatically without a user being present to authenticate to the system. Managed service accounts can be problematic in that they are often left in a default configuration, leaving them in roles or groups that have excessive privileges for what they need. Administrators frequently forget to periodically change service account passwords. These accounts tend to be "set and forget," and administrators often have no idea if the password meets complexity requirements or even the actual password.

When a managed service account is initially provisioned, normally when an application is installed or a system is built, you should review the account's permissions and roles to ensure that they are not by default far more extensive than it needs to perform its function.

You should place the managed service account in a specialized role or group that only has the needed privileges. Even so, confusing error messages may still occur that show that an application or process did not have sufficient privileges to execute a task; this is when you may have to adjust the privileges the account has, but at the same time you must be careful not to grant excessive privileges.

Account Access Review

As previously indicated, after initial provisioning has been performed as part of the onboarding process, a user's account and access rights must be managed for the duration of their employment. People get promoted, they transfer between departments, or they may face disciplinary actions such as suspension or investigation. In all these cases, the organization should make a point of reviewing the user's access rights and ensuring that they still meet the requirements of security clearance and need-to-know for their job function. Occasionally users must have

their privileges increased, be placed in new roles, or even have their access downgraded based on the requirements of the organization, particularly when users transfer and no longer need access to a particular system or information.

In addition to these event-based reviews, even if the user has no events that trigger a review during employment, the organization should routinely and periodically review a user's access rights to ensure that it is still current and valid. This could be as simple as reviewing user accounts whose last names start with a particular letter in the alphabet every week or reviewing users on the anniversary dates of their employment. In any case, there must be both event-based review and periodic review defined by policy. This ensures that users have not accumulated privileges they no longer need, that they are performing the job functions of and are assigned to roles that require access to systems and information, and that the provisioning and deprovisioning process is kept up to date.

You should review system accounts and service accounts on a scheduled, periodic basis. You should ensure that they are even still needed, since many applications may be uninstalled or upgraded and not require the service accounts any longer. You should also ensure that they still need the privileges granted to them and, if not, change those accordingly. You may also want to take this opportunity to change the service account password, since you don't want to leave it the same forever. Make sure that you do not make it the same password as other service accounts, because if one account is compromised, effectively all other accounts that share the same password will also be compromised. The password should be in accordance with organizational password complexity and length requirements as well.

Figure 5.5-1 summarizes the entire identity access management life cycle, including initial provisioning, managing access rights through the life of the account, and finally, deprovisioning an account. Note that this is just a generic life cycle; most organizations may have something similar but with a few variations.

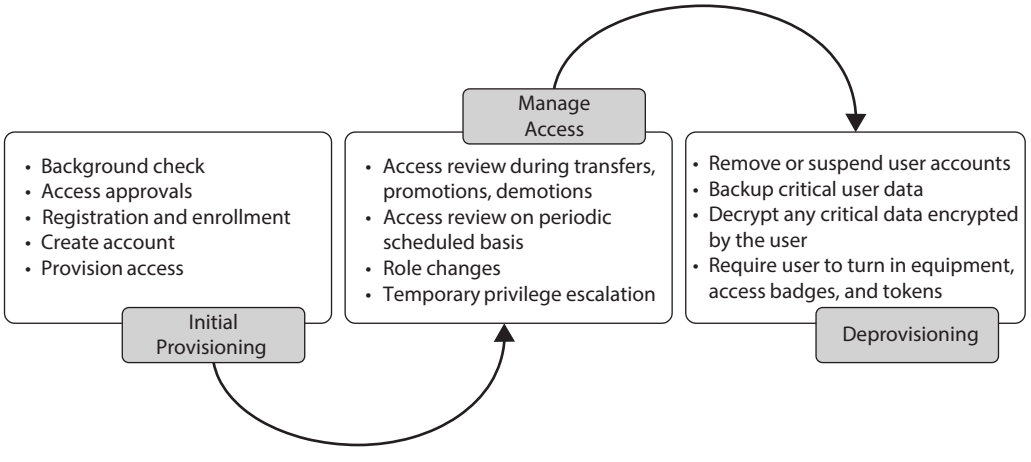


FIGURE 5.5-1 The identity access and management life cycle

REVIEW

Objective 5.5: Manage the identity and access provisioning lifecycle In this objective we reviewed the identity and access provisioning life cycle, which consists of all activities carried out to provide users with information systems access. This provisioning life cycle is based on the core security principles of identification and authentication, authorization, auditing and accountability, and nonrepudiation. The life cycle must also consider the principles of least privilege, separation of duties, and accountability.

We discussed the initial account provisioning for users, systems, and services, as well as the necessity to conduct periodic access review. We examined the provisioning and deprovisioning processes, which often require coordination between supervisors, human resources, and IT personnel to ensure that an individual has the right security clearance and need-to-know based on job function in order to be granted access to resources. We looked at how people are assigned to new roles so that access to systems and information can be carefully controlled by job function.

We also briefly touched on issues involving privilege escalation and the need to minimize additional privileges whenever possible. We looked at methods to temporarily escalate privileges, such as just-in-time provisioning, and the use of utilities such as `runas` and `sudo`. We examined the need to define privileges through roles when user accounts are initially provisioned, and then when a user is promoted or transferred.

Periodic account access reviews should be performed at various times while the user is employed. Account access should be reviewed when a user is promoted or transferred, or at least, as a minimum, on a periodic scheduled basis to ensure they have not accumulated excessive privileges and that they still need access to sensitive resources, based on their continuing need-to-know, job position, and security clearance.

5.5 QUESTIONS

1. You are onboarding a new employee and ensuring that they meet all security requirements prior to granting them an account and access to sensitive systems. One of the systems requires a rather extensive government background check, which has been delayed but is in progress. Which of the following is the best course of action for you to take in terms of granting access to systems and information?
   - **A.** Do not create the account or grant access to any other systems within the organization.
   - **B.** Create the account, and grant access only to those systems for which the individual is currently cleared and has a valid need-to-know due to job function, but no more.
   - **C.** Create the account and grant access to all systems the individual needs to access to perform their job, since the security clearance process is already underway and will likely come back as approved.
   - **D.** Create the account but only grant access to nonsensitive systems.

2. You are reviewing the privileges and access rights for an employee who has been with the company over ten years. During that time they have been promoted three times and transferred twice. When you review their access, you discover that years ago they were granted privileged access to systems that they no longer should have access to. Which of the following should you do?
   - **A.** Do nothing; the user may still need these access rights for continuity between positions.
   - **B.** Immediately take all the user's privileges away and suspend their account.
   - **C.** Continue to review all of their access rights, determine what rights they need for their current position, and take away all others.
   - **D.** Remove them from all roles and groups until your review is complete

5.5 ANSWERS

1. **B** If the new employee is otherwise approved for access to systems and information due to their qualifications and need-to-know, then you should create the account and grant them access only to those systems and information. You should not grant them access to any information or system for which they have not yet been approved, regardless of information sensitivity.
2. **C** You should continue to review their access, determine what access rights they need for their current position and rescind all other access rights. You should not take away any access they need for their current job, since they still need to perform their job duties. Doing nothing is not an acceptable answer since they would have access privileges they no longer require.



Objective 5.6 Implement authentication systems

This objective discusses authentication systems, particularly those that are connected to the Internet and help implement the concept of single sign-on between multiple systems. We will discuss the various authentication mechanisms and protocols that you need to be familiar with for the CISSP exam, which include OpenID Connect, OAuth, SAML, Kerberos, RADIUS, and TACACS+.

Authentication Systems

We have discussed authentication concepts throughout the book, but in this objective we take it further by discussing authentication mechanisms that must be connected together in a federated system to provide single sign-on (SSO) capabilities across different servers and resources

on the Internet. The authentication systems we will discuss support services that have become ubiquitous in the lives of everyone; for example, they enable you to log on to a financial site and securely pass on your credentials to your bank, pay bills to a different creditor, and transfer money. These protocols and mechanisms allow the different apps on smart devices to use identity management (IdM) services for pass-through authentication in a secure manner. We will also discuss more traditional authentication mechanisms that you will also encounter, such as Kerberos as it is implemented in Microsoft's Active Directory, and the remote access protocols RADIUS and TACACS+.

Open Authorization

Open Authorization (OAuth) is an open-standards authorization framework. If you've ever used an app on your smartphone that required access to sensitive information or services from another app, you may have been prompted with a requirement to authenticate yourself, in the form of a pop-up box that asks for your user ID and password for the app that must provide the services. This is an example of OAuth in practice. OAuth exchanges messages between the application programming interfaces (APIs) of applications and creates a temporary token showing that access to the information or services provided by one application to the other is authorized. Note that OAuth is used for authorization only, not authentication. Authentication is where OpenID and OpenID Connect come in, discussed next.

OpenID Connect

OpenID Connect (OIDC) is another open standard, but in this case, it is an *authentication* standard supported by the OpenID Foundation. OpenID can facilitate a user logging into several different websites using only one set of credentials. The set of credentials is maintained by a third-party identity provider (IdP), called an OpenID provider. If a user visits a web resource that supports OpenID, they must enter their OpenID credentials, which in turn causes the website (called a *relying party*) and the OpenID provider to exchange authentication information, and the access is allowed if the credentials and the user's identity are authenticated by the provider.



EXAM TIP To reduce the confusion, remember that OAuth is an *authorization* framework, OpenID is an *authentication* layer, and OpenID Connect (OIDC) is an authentication framework that sits on top of OAuth to provide authentication and authorization services.

Security Assertion Markup Language

The Security Assertion Markup Language (SAML) is an XML-based language used to allow identity providers to securely exchange authentication information, such as identities and attributes, with service providers. SAML uses XML-based messages, called *SAML assertions*,

that are exchanged between a user requesting access to a resource and the owner of the resource. These assertions contain credential information, such as identity, which is authenticated by a centralized third-party IdP.

SAML uses specific XML tags to allow applications to format identification information about the user. There are generally three components to SAML, which is now in version 2.0:

- **The "principal"** This is the user requesting access.
- **The identity provider** This is the entity that authenticates the user's identity.
- **The service provider** This is the entity who owns or controls the service or application to which the user is requesting access.

Since SAML is a common standard, it is used across many different identity and service providers. If the user requests access to a website or resource, for instance, the service provider requests identification from the user, who submits it through a web-based application. The service provider takes that response and requests authentication verification from the identity provider, who in turn provides the service provider with validated authentication information, such as assertions regarding the user's authentication information and any authorizations they have for the resource. The service provider can then allow or deny access to the resource. Examples of third-party IdPs that can authenticate users to web resources include Microsoft, Google, and Facebook.

Kerberos

Kerberos is a popular open-standards authentication protocol that is used in a wide variety of technologies, including Linux and UNIX, and most prominent among them is Microsoft's Active Directory. Kerberos was developed at MIT and provides SSO and end-to-end secure authentication services. It works by generating authentication tokens, called *tickets*, which are securely issued to users so they can further access resources on the network. Kerberos uses symmetric key cryptography and shared secret keys, rather than transmitting passwords over the network.

Kerberos has several important components that you should be aware of for the CISSP exam:

- The *Key Distribution Center (KDC)* is the primary component within a Kerberos realm. It stores keys for all of the principals and provides authentication and key distribution services. In Microsoft Active Directory, a domain controller serves as the KDC.
- A *principal* is any entity that the KDC has an account for and shares a secret key with.
- *Tickets* are essentially session keys which expire after a predetermined amount of time. Tickets are issued by the KDC and are used by principals to access resources.
- A *ticket granting ticket (TGT)* is issued to the user when they authenticate so that they can use it to later request session tickets.
- The Kerberos *realm* consists of all the security principals for which the KDC provides services; Microsoft implements a Kerberos realm as an Active Directory domain.

- The *Authentication Service (AS)* on the KDC authenticates principals when they log on to the Kerberos realm.
- The *ticket granting service (TGS)* generates session tickets that are provided to security principals when they need to authenticate to a resource or each other.

The process for authenticating using Kerberos can be a complex one, but you should be familiar with it for the exam. Essentially the process works like this:

1. A user authenticates by logging into a domain host, which sends the username to the authentication service on the KDC. The KDC sends the user a TGT, which is encrypted with the secret key of the TGS.
2. If the user has correctly entered their password upon login, the TGT is decrypted, and the user is allowed access to their device.
3. When the user needs access to a network resource, their workstation sends the TGT to the TGS to request access to the resource.
4. The TGS creates and sends an encrypted session ticket back to the user to authenticate to the resource.
5. When the user receives the encrypted session ticket, their workstation decrypts it and sends it on to the resource for authentication.
6. If the resource (e.g., a file server or printer) is able to authenticate that the ticket came from both the user and the KDC, then access is allowed.

Note that authentication and ticket distribution are highly dependent upon system time synchronization in a Kerberos realm; the KDC timestamps all tickets and must have access to a valid time source, such as a Network Time Protocol (NTP) server. If the time settings on devices within the Kerberos realm vary by more than a specified amount, it will create issues with the timestamps on the tickets. This synchronization requirement helps prevent replay attacks.

Kerberos is a very resilient system for authentication and authorization; however, it also has several weaknesses:

- The KDC is a single point of failure unless there is redundancy in place with multiple servers.
- The KDC must be scalable and able to handle the number of users on the network.
- Weak keys, such as passwords, are vulnerable to attack.
- Secret keys are temporarily stored on user devices (sometimes in a decrypted state), so if the device is compromised, the keys may also be compromised.

The operating system and other security controls can help mitigate Kerberos weaknesses; these include encrypting all traffic and end-user device protection.

Remote Access Authentication and Authorization

The CISSP exam objectives require you to understand various remote access authentication and authorization protocols and technologies. We covered some of these earlier in the domain and briefly mentioned both RADIUS and TACACS+ in Objective 4.3. For this objective we will go into more detail regarding RADIUS, TACACS+, and another remote access protocol, Diameter. These protocols are referred to as authentication, authorization, and auditing (or accounting, depending on the source), or AAA, protocols.

AAA protocols provide these services to remote users who may use dial-up, VPN, or other types of remote connections. They make use of remote access servers, running these protocols, that can also connect to authentication databases, such as Microsoft's Active Directory. Although these remote access technologies are not widely used, having been replaced by more modern VPN technologies, you may still see them used in special circumstances (such as an initial connection to an Internet provider prior to full authentication and connection provisioning), legacy systems, or locations where there is sparse high-speed Internet availability.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is an access protocol that can provide authentication, authorization, and accounting services for remote connections into networks. RADIUS was originally created to support dial-up connections into larger Internet service providers (ISPs) using infrastructure with large modem banks, authentication databases, and accounting and billing services for those connections. RADIUS uses a network access server for dial-in but can then offload authentication services to internal credential databases. RADIUS supports many of the older authentication protocols, such as PAP, CHAP, and MS-CHAP.

In terms of communications protection, the RADIUS server can encrypt passwords but nothing else, including username and other session communications information, such as IP address, hostname, and so on. RADIUS also requires more overhead than other protocols, particularly TACACS+, discussed in the next section, since it relies on UDP (over ports 1812/1813 or 1645/1646) and must have built-in error checking mechanisms to make up for UDP's connectionless design.

TACACS+

Terminal Access Controller Access Control System Plus (TACACS+) functions like RADIUS but is very different. Whereas RADIUS uses UDP, TACACS+ uses TCP (over port 49). There are three distinct protocols in the TACACS family: the original TACACS protocol, TACACS+, and Extended TACACS (XTACACS). You should be aware that the original TACACS, XTACACS, and TACACS+ are three different protocols; they're not compatible with each other. In fact, TACACS+ and XTACACS started out as Cisco proprietary protocols, but TACACS+ has since become an open standard.

TACACS combines its authentication and authorization services, but XTACACS separates the authentication, authorization, and accounting functions. It's also worthy to note that TACACS+ improves upon XTACACS by adding multifactor authentication, whereas the original TACACS uses only fixed passwords.

Diameter

Diameter was developed to enhance the functionality of RADIUS and overcome some of its many limitations. It provides more flexibility and capabilities than either RADIUS or TACACS+. It is used by wireless devices and smartphones, and can be used over Mobile IP, Ethernet over PPP, and even VoIP. Diameter is different than both RADIUS and TACACS+ in that while they are client/server protocols, Diameter is a peer-to-peer protocol where either endpoint can initiate communications. Note that Diameter is not backward compatible with RADIUS; it is considered an upgrade.



EXAM TIP Although not specifically called out in the CISSP exam objectives, you may see Diameter as a distractor or incorrect answer when a question focuses on RADIUS or TACACS+, so it's a good idea to at least be familiar with it. Diameter is a similar technology and is almost always discussed alongside RADIUS and TACACS+ protocols.

REVIEW

Objective 5.6: Implement authentication systems In this objective we discussed authentication mechanisms, systems, and protocols that are widely used to connect systems on the Internet and securely pass credentials from one system to another, enabling single sign-on for users. We discussed the particulars of two important authentication mechanisms, Open Authorization (OAuth) and OpenID Connect (OIDC), as well as an XML-based language, SAML, that allows credentials to be securely exchanged between systems that depend on third-party IdM services. We also discussed protocols such as Kerberos, which is prevalent in Microsoft's Active Directory, and three key AAA protocols used in remote connections for authentication, authorization, and accounting, RADIUS, TACACS+, and Diameter.

5.6 QUESTIONS

1. You are performing a security review of a web-based application developed by your company. However, the developers have not built in support for any third-party identity provider and are unsure of how to do so. They need to be able to use a common language to help exchange authentication information with third-party providers. Which of the following should you advise them to use?
   - **A.** OAuth
   - **B.** RADIUS
   - **C.** SAML
   - **D.** TACACS+

2. You are troubleshooting a client's Active Directory infrastructure. They are having problems with users only sporadically being able to authenticate to the network and access resources. After looking at multiple logs, you see several entries indicating that tickets are expiring rapidly or are being rejected because they are no longer valid. Which of the following could be the source of the problem?
   - **A.** Time synchronization across the network is not working.
   - **B.** The KDC is offline.
   - **C.** Users are inputting incorrect credentials.
   - **D.** Faulty network equipment is causing a disruption of service.

5.6 ANSWERS

1. **C** You should advise the developers that the web application should support SAML, since it is an XML-based standardized language used for exchanging authentication information with third-party identity providers.
2. **A** Since tickets are expiring before they can be used, and some logs indicate that tickets are being rejected because they are no longer valid, you should start investigating the network time source and ensure that all hosts are receiving correct time synchronization. Kerberos is highly dependent on time synchronization throughout the network, and incorrect timing could result in tickets expiring before they should, or tickets being rejected because they are no longer valid. This would also sporadically prevent users from authenticating to the network or accessing resources.



Security Assessment and Testing

Domain Objectives

- **6.1** Design and validate assessment, test, and audit strategies.
- **6.2** Conduct security control testing.
- **6.3** Collect security process data (e.g., technical and administrative).
- **6.4** Analyze test output and generate report.
- **6.5** Conduct or facilitate security audits.

Domain 6 addresses the important topic of security assessment and testing. Security assessment and testing enable cybersecurity professionals to determine if security controls implemented to protect assets are functioning properly and to the standards they were designed to meet. We will discuss several types of assessments and tests during the course of this domain, and address situations and reasons why we would use one over another. This domain covers designing and validating assessment, test, and audit strategies; conducting security control testing; collecting the security process data used to carry out a test; examining test output and generating test reports; and conducting security audits.



Objective 6.1 Design and validate assessment, test, and audit strategies

Before conducting any type of security assessment, test, or audit, the organization must pinpoint the reason it wants to conduct the activity in order to select the right process to meet its needs. Information security professionals conduct different types of assessments, tests, and audits for various reasons, with distinctive goals, and sometimes use different approaches for each. In this objective we will examine the reasons why an organization may want to conduct specific types of assessments, tests, and audits, including internal, external, and third-party, and discuss the strategies for each.

Defining Assessments, Tests, and Audits

Before we discuss the reasons for, and the strategies related to, conducting assessments, tests, and audits, it's useful to define what each of these processes means. Although the terms assessment, test, and audit are frequently used interchangeably, each has a distinctive meaning. Note that each of these security processes consists of activities that can overlap; in fact, any or all three of these processes can be conducted simultaneously as part of the same effort, so it can sometimes be difficult to differentiate one from another or to separate them as distinct events.



NOTE Although "assessment" is the term that you will hear most often even when referring to tests and audits, for the purposes of this objective we will use the term "evaluation" to avoid any confusion when generally referencing security assessments, tests, and audits.

A *test* is a defined procedure that records properties, characteristics, and behaviors of the system being tested. Tests can be conducted using manual checklists, automated software such as vulnerability scanners, or a combination of both. Typically, there is a specific goal of the test, and the data collected is used to determine if the objective is met. The test results may

be compared to regulatory standards or compliance frameworks to see if those results meet or exceed the requirements. Most security-related tests are technical in nature, but they don't necessarily have to be. For example, a social engineering test can help determine if users are sufficiently trained to identify and resist social engineering attempts. Other examples of tests that we will discuss later in the domain include vulnerability tests and penetration tests.

An *assessment* is a collection of related tests, usually designed to determine if a system meets specified security standards. The collection of tests does not have to be all technical; a control assessment, for example, may consist of technical testing, documentation reviews, interviews with key personnel, and so on. Other assessments that we will discuss later in the domain include vulnerability assessments and compliance assessments.

Cross-Reference

The realm of security assessments also includes risk assessments, which were discussed in Objective 1.10.

Audits are conducted to measure the effectiveness of security controls. Auditing as a routine business process is conducted on a continual basis. Auditing as a detective control usually involves looking at audit trails such as log files to discover wrongdoing or violations of security policy. On the other hand, an *audit* as a specific event consists of a tailored, systematic assessment of a particular system or process to determine if it meets a particular standard. While audits can be used to analyze specific occurrences within the infrastructure, such as transactions, auditors may be specifically tasked with reviewing a particular incident, event, or process.

Designing and Validating Evaluations

Evaluations are designed to achieve specific results, whether it is a compliance assessment, a security process audit, or a system security test. While most of these types of evaluations use similar processes and procedures, the information and results that the organization needs to gain from the evaluation may affect how it is designed. For example, a test that must yield technical configuration findings must be designed to scan a system, collect security configuration information in a commonly used format, and report those results in a consistent manner that has been validated as accurate and acceptable. Design and validation of evaluations are critical to consistency and the usefulness of the results.

Goals and Strategies

The type of security evaluation an organization chooses to use depends on its goals. What does the organization want to get out of the assessment? What does it need to test? How does it need to conduct an audit and why do they need to conduct it in a specific way? Depending on what type of security evaluation the organization uses, there are different strategies that can be used to conduct each type of activity. For the purposes of covering exam objective 6.1, we will

discuss goals and strategies that involve the use of internal personnel, external personnel, and third-party teams for assessments, tests, and audits.

Use of Internal, External, and Third-Party Assessors

Any type of evaluation can be resource intensive; a simple system test, even using internal personnel, costs time and labor hours. A larger assessment that covers multiple systems or an audit that looks at several different security business processes can be very expensive, especially when it involves the use of many different personnel and requires system downtime. An organization may be able to reduce cost by making use of available internal resources in an effective and efficient way, especially since the labor hours of those resources are already budgeted. However, internal teams are not always the best choice to perform an evaluation, as they can have confirmation biases. Sometimes the use of external teams or third-party assessors is necessary to get a truly objective evaluation.

Table 6.1-1 summarizes the use of internal assessors versus external (second party) assessors and third-party assessors, as well as their advantages and disadvantages.

TABLE 6.1-1 Types of Assessors and Their Characteristics

Type of Assessor	Description	Advantages	Disadvantages
Internal	The assessors are already employed by the organization as IT or cybersecurity personnel.	Less expensive; more familiar with the infrastructure.	Inadequate or limited training and evaluation tools; limited exposure to new evaluation techniques; potential for conflicts of interest and exposure to organizational politics and influence.
External	The assessors work for a business partner or on behalf of a business partner as part of contract requirements.	Expenses may be billed to the contract with the business partner.	Less familiar with the organization infrastructure; evaluation includes only very specific areas as required by contract.
Third-party	The assessors are an external team independent of the organization or its business partners, hired specifically to conduct specific evaluations and audits.	Immune to internal organizational conflicts and politics; access to a wide variety of tools and assessment techniques; usually more experienced than an internal team.	More expensive option; less familiar with the organizational infrastructure; may gain access to sensitive knowledge, requiring a nondisclosure agreement.



EXAM TIP Although the difference between internal and other types of assessors seems simple, the difference between external and third-party assessors may be less so. Internal assessors work for the organization; external assessors work for a business partner or someone connected to the organization. Third-party assessors are completely independent of the organization and any partners. Typically, they work for an independent outside entity such as a regulatory agency.

REVIEW

Objective 6.1: Design and validate assessment, test, and audit strategies In this objective we defined and discussed assessments, tests, and audits. We also looked at three key strategies for the use of internal, external, and third-party assessors. The choice of which of these three types of assessors to use depends on several factors: the organization's security budget, the potential for conflicts of interest or politics if internal auditors are used, and assessor familiarity with the infrastructure. An additional consideration is a broader knowledge of assessment techniques and access to more advanced tools.

6.1 QUESTIONS

1. Your manager has tasked you to determine if a specific control for a system is functional and effective. This effort will be very limited in scope and focus only on specific objectives using a defined set of procedures. Which of the following would be the most appropriate method to use?
   - **A.** Assessment
   - **B.** Test
   - **C.** Audit
   - **D.** Review
2. You have a major security compliance assessment that must be completed on one of your company's systems. There is a very large budget for this assessment, and it must be conducted independently of organizational influence. Which of the following types of assessment teams would be the most appropriate to conduct the assessment?
   - **A.** Third-party assessment team
   - **B.** Internal assessment team
   - **C.** External assessment team
   - **D.** Second-party assessment team

6.1 ANSWERS

1. **B** A test would be the most appropriate method to use, since it is very limited in scope and involves only the list of procedures used to check for the control's functionality and effectiveness. An assessment consists of several tests and is much broader in scope. An audit focuses on a specific event or business process. A review is not one of the three types of assessments.
2. **A** A third-party assessment team would be the most appropriate to conduct a major security compliance assessment. This type of assessment team is independent from the organization's influence, but it's usually also the most expensive type of team to employ for this effort. However, the organization has a large budget for this effort. Internal assessment teams are not independent of organizational influence but are usually less expensive. "External assessment team" and "second-party assessment team" are two terms for the same type of team, usually employed as part of a contract between two entities, such as business partners, and used only to assess specific requirements of the contract.



Objective 6.2 Conduct security control testing

Objective 6.2 delves into the details of security control testing using various techniques and test methods. Security control testing involves testing a control for three main reasons: to evaluate the effectiveness of the control as implemented, to determine if the control is compliant with governance, and to see how well the control reduces or mitigates risk for the asset.

Security Control Testing

In this objective we will discuss various tests that serve to verify that security controls are functioning the way they were designed to, as well as validate that they are effective in fulfilling their purpose. Each of these tests serves to verify and validate security controls. We will examine the different types of testing, such as vulnerability assessment and penetration testing. We will also take a look at reviewing audit trails in logs, which describe what happened during a test. We will review other aspects of security control testing, such as synthetic transactions used during testing, as well as some of the more specific types of tests we can conduct on infrastructure assets. These include code review and testing, use and misuse case testing, interface testing, simulating a breach attack, and compliance testing. During this discussion, we will also address test coverage analysis and exactly how much testing is enough for a given assessment.

Vulnerability Assessment

A vulnerability assessment is a type of test that involves identifying the vulnerabilities, or weaknesses, present in an individual component of a system, an entire system, or even in different areas of the organization. While we most often look at vulnerability assessments as technical types of tests, vulnerability assessments can be performed on people and processes as well. Recall from Objective 1.10 that a vulnerability is not only a weakness inherent to an asset, but also the absence of a security control, or a security control that is not very effective in protecting an asset or reducing risk.

Vulnerability assessments can be performed using a variety of methods, but most technical assessments involve the use of network scanning tools, such as Nmap, OpenVAS, Qualys, or Nessus. These vulnerability scanners give us information about the infrastructure, such as live hosts on the network, open ports, protocols and services on those hosts, misconfigurations, missing patches, and a wealth of other information that can, when combined with other types of tests, inform us as to what weaknesses and vulnerabilities are present on our network.



EXAM TIP Vulnerability assessments are distinguished from penetration tests, discussed next, in that during a vulnerability assessment we only look for vulnerabilities; we don't act on them by attempting to exploit them.

Vulnerability mitigation is also covered in other areas in this book, but understand that you must prioritize vulnerabilities you discover for mitigation. Most often mitigations included patching, reconfiguring devices or applications, or implementing additional security controls in the infrastructure.

Penetration Testing

A penetration test can be thought of as the next logical extension of a vulnerability assessment, with some key differences. First, while a vulnerability assessment looks for weaknesses that theoretically can be exploited, a penetration test takes it one step further by proving that those weaknesses *can* be exploited. It's one thing to discover a vulnerability that a scanning tool reports as being severe, but it's quite another to have the tools and techniques or the knowledge necessary to actually exploit the weakness. In other words, there's often a big difference between discovering a severe vulnerability and the ability of an attacker to take advantage of it. Penetration testing attempts to exploit discovered vulnerabilities and affect system changes. Penetration tests are more precise in that they demonstrate the true weaknesses you should be concerned about on the network. This helps you to better apply resources and mitigations in an effort to eliminate those vulnerabilities that can truly be exploited, versus the ones that may not easily be exploitable.



CAUTION Penetration tests can be very intrusive on an organization and can result in the failure or degradation of infrastructure components, such as systems, networks, and applications. Before performing a penetration test, all parties should be aware of the risks involved, and there should be a documented agreement between the penetration test team and the client that explicitly states so, with the client signing that they understand the potential risks.

There are several different ways to categorize both penetration testers and penetration tests. First, you should be aware of the different types of testers (sometimes referred to as hackers) normally involved in penetration tests:

- **White-hat (ethical) hackers** Security professionals who test a system to determine its weaknesses so they can be mitigated and the system better secured
- **Black-hat (unethical) hackers** Malicious entities who hack systems to infiltrate them, access sensitive data, or disrupt infrastructure operations
- **Gray-hat hackers** Hackers who go back and forth between the white-hat and black-hat worlds, sometimes selling their expertise for the good of an organization, and sometimes engaging in unethical or illegal activities



NOTE Black-hat hackers (and gray-hat hackers operating as black-hat hackers) aren't really considered penetration testers, as their motives are not to test security controls to improve system security.

Note that ethical hackers or testers may be internal security professionals who work for the organization on a continual basis or they may be external security professionals who provide valuable services as an independent team contracted by organizations.

In addition to the various colors of hats we have in penetration testing, there are also red teams, blue teams, and white cells, which we encounter during the course of a security test or exercise. The term red team is merely another name for the penetration testing team, the attacking group of testers. A blue team is the name of the group of people who serve as the defenders of the infrastructure and are tasked with detecting and responding to an attack. Finally, the *white cell* operates as an independent entity facilitating communication and coordination between the blue and red teams and organizational management. The white cell is usually the team managing the exercise, serving as liaison to all stakeholders, and having final decision-making authority over any conflicts during the exercise.

Just as there are different categories of testers, there are also different categories of tests. Each category has its own distinctive characteristics and advantages. They are summarized here as key terms and definitions to help you remember them for the exam:

- **Full-knowledge (aka *white box*) test** A penetration test in which the test team has full knowledge about the infrastructure and how it is architected, including operating systems, network segmentation, devices, and their vulnerabilities. This type of test is useful for enabling the team to focus on specific areas of interest or particular vulnerabilities.
- **Zero-knowledge (aka *blind* or *black box*) test** A penetration test in which the team has no knowledge of the infrastructure and must rely on any open source intelligence it can discover to determine network characteristics. This type of test is useful in discovering the network architecture and its vulnerabilities from an attacker's point of view, since this is the most likely circumstances a real attacker will face. In this scenario, the defenders may be aware that the infrastructure is being attacked by a test team and will react accordingly per the test scenarios.
- **Partial-knowledge (aka *gray box*) test** A penetration test that takes place somewhere between full- and zero-knowledge testing, where the test team has only some limited useful knowledge about the infrastructure.
- **Double-blind test** A penetration test that is zero-knowledge for the attacking team, but also one in which the defenders are not aware of an assessment. The advantage to this type of test is that the defenders are also assessed on their detection and response capabilities, and tend to react in a more realistic manner.
- **Targeted test** A test that focuses on specific areas of interest by organizational management, and may be carried out by any of the mentioned types of testers or teams.



EXAM TIP You should never carry out a penetration test unless you are properly authorized to do so. Organizational management and the test team should jointly develop a set of "rules of engagement" that both clearly defines the test parameters and grants you permission to perform the test in writing.

Log Reviews

Log reviews serve several functions, both during normal security business processes and during security assessments. On a broader scope, log reviews take place on a daily basis to inspect the different types of transactions that occur within the network, such as auditing user actions or other events of interest. Logs are reviewed on a frequent basis to determine if any anomalous

or malicious behavior has occurred. During a security assessment, log reviews serve the same function and also record the results of security testing. Logs are reviewed both during and after security tests to ensure that the test happened according to its designed parameters and produced predictable results. When test results occur that were not predicted, logs can be useful in tracing the reason.

Logs can be manually reviewed, or ingested by automated systems, such as security information and event management (SIEM) systems, for aggregation, correlation, and analysis. Logs are useful in reconstructing timelines and an order of events, particularly when they are combined from different sources that may present various perspectives.

Synthetic Transactions

Information systems normally operate on a system of transactions. A *transaction* is a collection of individual events that occur in some sort of serial or parallel sequence to produce a desired behavior. A transaction that occurs when a user account is created, for example, consists of several individual events, such as creating the username, then its password, assigning other attributes to the account, and finally adding the account to a group of users that have specific privileges. *Synthetic transactions* are automated or scripted events that are designed to simulate the behaviors of real users and processes. They allow security professionals to systematically test how critical security services behave under certain conditions. A synthetic transaction could consist of several scripted events that create objects with specific properties and execute code against those objects to cause specific behaviors. Synthetic transactions have the advantage of being controlled, predictable, and repeatable.

Code Review and Testing

Code review consists of examining developed software code's functional, performance, and security characteristics. Code reviews are designed to determine if code meets specific standards set by the organization. There are several methods for code review, including manual reviews conducted by people and automated reviews conducted by code review software and scripts. Think of code review as proofreading software code before it is published and put into production in operational systems. Some of the items reviewed in code include structure and syntax, variables, interaction with system resources such as memory and CPU, security issues such as embedded passwords and encryption modules, and so on.

In addition to review, code is also tested using both manual and automated tools. Some of the security aspects of code testing involve

- Input validation
- Secure data storage
- Encryption and authentication mechanisms

- Secure transmission
- Reliance on unsecured or unknown resources, such as library files
- Interaction with system resources such as memory and CPU
- Bounds checking
- Error conditions resulting in a nonsecure application or system state

Misuse Case Testing

Use case testing is employed to assess all the different possible scenarios under which software may be correctly used from a functional perspective. *Misuse* case testing involves assessing the different scenarios where software may be misused or abused. A misuse case is a type of use case that tests how software code can be interacted with in an unauthorized, malicious, or destructive way. Misuse cases are used to determine all the various ways that a malicious entity, such as a hacker, could exploit weaknesses in a software application or system.

Test Coverage Analysis

When performing an assessment, you should determine which parts of a system, application, or process will be covered by the assessment, how much you're going to test, and to what depth or level of detail. This is referred to as *test coverage analysis* and is performed to derive a percentage measurement of how much of a system is going to be examined by a test or overall assessment. This limit is something that IT professionals, security professionals, and management want to know before approving and conducting an assessment.

Test coverage analysis considers many factors, such as the cost of the assessment, personnel availability, and system criticality. Obviously, highly critical systems require more test coverage than noncritical systems, to ensure that all possible scenarios and exceptions are analyzed. Test coverage analysis can be one of those measurements that can be extended over time; for example, an organization could conduct in-depth vulnerability testing on 25 percent of its assets each week so that after a month, all assets have been covered, and then the process repeats. This approach ensures that no one piece of the infrastructure is left untested.

Interface Testing

Interfaces are connections and entry/exit points between hardware and software. Examples of interfaces include a graphical user interface (GUI) that helps a human user communicate with a system and a network interface that connects systems to networks to facilitate the exchange of data. Interfaces can also include security mechanisms, application programming interfaces (APIs), database tables, and a multitude of other examples. In any event, interfaces represent critical junction points for communications and data exchange and should be tested for security.

Security issues with interfaces include data movement from a more secure environment to a less secure one, introduction of malware into a system, and unauthorized access by a user or process. Interfaces are often responsible for not only exchanging data between systems or networks, but also data transformation, which can affect data integrity. Interface testing should address all of these issues and ensure that data exchanged between entities is of the right format and transferred with the correct access controls.

Breach Attack Simulations

Breach attack simulations are a form of automated penetration test that the organization may conduct on itself, targeting various systems, network segments, or even applications. A breach attack simulation is normally done on a regular schedule and lies somewhere along the spectrum between vulnerability testing and penetration testing, since it attempts to find vulnerabilities that are in fact exploitable. Normally, a breach attack simulation will not actually exploit a vulnerability, unless it is purposely configured to do so; it may run automated tests and synthetic transactions to show that the vulnerability was, in fact, exploitable, without actually executing the attack. An advantage of a breach attack simulation is that since it's performed on a regular basis, it's not simply a point-in-time assessment, as are penetration tests. Breach attack simulations often come as part of a software package, sometimes as part of a subscription- or cloud-based security suite.

Compliance Checks

Compliance checks are tests or assessments designed to determine if a control or security process complies with governance requirements. While compliance checks are conducted in pretty much the same way as the other types of assessments discussed so far, the real difference is that the results are checked against sometimes detailed requirements passed down through laws, regulations, or mandatory standards that the control must meet. For example, it may not be enough to confirm that the control ensures that data is encrypted during transmission; the governance requirement may mandate that the encryption strength be at a certain level or use a specific algorithm, so the control must be checked to determine if it meets that requirement.



NOTE While definitely interconnected, the terms *secure* and *compliant* are not synonymous. A control could be secure, but not necessarily compliant due to a lack of documentation or a consistent process. The opposite is also true.

Most assessments examine not only the security of a control and how well it protects an asset, but also if it complies with prescribed governance. An individual test, such as a vulnerability assessment performed by a network scanner, may only determine if the control is secure. It's up to the security analyst performing the overall assessment to determine its compliance with governance requirements.

REVIEW

Objective 6.2: Conduct security control testing In this objective we looked a bit more in depth at security control testing. We discussed overall assessments, such as vulnerability assessments and penetration testing. Vulnerability assessments look for weaknesses in systems but do not attempt to exploit them. Penetration testing, on the other hand, not only finds vulnerabilities but assesses their severity by attempting to exploit them. Penetration testing can come in different flavors and can be performed by different types of testers, such as ethical hackers or internal and external testers. Penetration testing can be categorized in different ways, such as full-knowledge, zero-knowledge, and partial-knowledge tests.

Different tools and methods can be used for assessments but should include various methods to verify and validate security controls. Verification means that the control is tested to determine if it is working as designed; validation means that the control's actual effectiveness in performing that function is determined. Security control testing includes:

- Log reviews to determine if the results of the tests are consistent with what we expect.
- Synthetic transactions, which are scripted sets of events that can be used to test different aspects of functionality or security with a system.
- Code review and testing, which are critical to ensuring that software code is error-free and employs strong security mechanisms.
- Misuse case testing, which looks at how users might abuse systems, resulting in unauthorized access to information or system compromise.
- Test coverage analysis, which examines data based on a percentage of measurement of how much of a system or application is tested.
- Interface testing, which examines the different connections and data exchange mechanisms between systems, networks, applications, and other components.
- Breach attack simulation, which is an automated periodic process that not only finds vulnerabilities but more closely examines whether or not they can be exploited, without actually doing so.
- Compliance checks, which ensure that controls are not only secure but also meet governance requirements.

6.2 QUESTIONS

1. You are a cybersecurity analyst who has been tasked with determining what weaknesses are on the network. Your supervisor has specifically said that regardless of the weaknesses you find, you must not disrupt operations by determining if they can be exploited. Which of the following is the best type of assessment you could perform that meets these requirements?
   A. Penetration test
   B. Vulnerability assessment

C. Zero-knowledge test

D. Compliance check

2. Your organization has performed several point-in-time tests to determine what weaknesses are in the infrastructure and if they can be exploited. However, these types of tests are expensive and require a lot of planning to execute. Which of the following would be a better way to determine on a regular basis not only weaknesses but if they could be exploited?

   A. Perform a vulnerability assessment

   B. Perform a zero-knowledge penetration test

   C. Perform a full-knowledge penetration test

   D. Perform a breach attack simulation

6.2 ANSWERS

1. **B** Of the choices given, performing a vulnerability assessment is the preferred type of test, since it discovers vulnerabilities in the infrastructure without disrupting operations by attempting to exploit them.

2. **D** A breach attack simulation is an automated method of periodically scanning and testing systems for vulnerabilities, as well as running scripts and other automated methods to determine if those vulnerabilities can be exploited. A breach attack simulation does not actually perform the exploitation on vulnerabilities, however.



Objective 6.3 Collect security process data (e.g., technical and administrative)

To address objective 6.3, we will discuss the importance of collecting and using security process data for purposes of security assessment and testing. *Security process data* is any type of data generated that is relevant to managing the security program, whether it is routine data resulting from normal security processes or data that comes from specific events, such as incidents or test events. In the context of this domain, we will focus on data generated as a result of security business processes.

Security Data

There are so many sources of data that can be collected and used during security processes that they could not possibly be covered in this limited space. In general, you will want to collect both technical and administrative process data, whether it is electronic data, written

data, or narrative data. The various data collection methods include ingesting data from log files, configuration files, and so on, as well as manual documentation reviews and interviews with key personnel.



EXAM TIP Security process data can be technical or administrative. Administrative data consists of analyzed or narrative data from administrative processes (such as metrics, for example), and technical data comes from technical tools used to collect data from systems.

Security Process Data

Although there are some generally accepted categories of information that we use for security, any data could be "security data" if it supports the organization's security function. Obviously, security data includes logs, configuration files, and audit trail data, but other types of data, such as security process data (data that results from an administrative or technical process), risk reports, and other information can be useful in refining security processes as well. Beneficial data that is not specifically created for security purposes includes architectural, electric circuit, and building diagrams, as well as videos or voice recordings and information collected for unrelated objectives. Some of this information may be used to create a timeline for an incident during an investigation, for example, or to support a security decision. In the next few paragraphs, we will cover different sources of security process data and how they apply to the security function.

Data Sources

Data may be received and collected from a wide variety of sources within the infrastructure through automated means or manual collection. Technical data sources can be agent-based software on devices, or data can be collected by running a program or script that retrieves it from a host manually. Indeed, you should look at any source that will provide information regarding the effectiveness of a security control, the severity of a vulnerability, or an event that should be investigated and monitored. Valuable information may be gathered from a variety of sources, including electronic log files, paper access logs, vulnerability scan results, and configuration data. In a mature organization, a centralized security information and event management (SIEM) system may collect, aggregate, and correlate all of this information.

Cross-Reference

Objective 7.2 covers SIEM in more depth.

Account Management

Account management data is of critical importance to security personnel. Account management data includes details on user accounts, including their rights, permissions and other privileges, and other user attributes. Account usage history usually comes in electronic form as an account listing, but it could also come from paper records that are signed by supervisors to grant accounts to new individuals or increase their privileges. This information can be used to match user identities with audit trails that facilitates auditing and accountability.

Backup Verification Data

Backup verification data is important in that the organization needs to be sure that its information is being properly backed up, secured, transported, and stored. Backup verification data can come either from a written log in which a cybersecurity analyst or IT professional has manually recorded that backups have taken place or, more commonly, from transaction logs produced by the backup application or system. All critical information should be backed up in the event that an incident occurs that renders data unusable or damages a system. The event that causes a data loss should be well documented, along with the complete restore process for the backed-up data.

Training and Awareness

Since security training is the most critical piece of managing the human element, individuals should have training records that reflect all forms of security training completed. When the organization understands how well its personnel are trained, how often, and on which topics, this data can later be compared to trends that indicate whether training is responsible for the number or severity of security incidents increasing or decreasing. Training data can be recorded electronically in a database so effective training can be analyzed based on its frequency, delivery method, and/or facilitator. Security training data is a good indicator for how well personnel understand their security responsibilities.

Disaster Recovery and Business Continuity Data

The ability to restore data from backups is important. However, the bigger picture is how well the organization reacts to and recovers from a negative event, such as an incident or disaster. Data regarding how quickly a team responds to an incident or how well the team is organized and equipped is essential to assessing readiness to meet recovery objectives. Disaster recovery (DR) and business continuity (BC) data provide details of critical recovery point objectives, recovery time objectives, and maximum allowable downtime. The organization should develop these metrics before a disaster occurs. The most important data is how well the organization actually meets these objectives during an event.

Cross-Reference

Refer to Objective 1.8 for coverage of BC requirements and to Objective 7.11 for coverage of DR processes.

Key Performance and Risk Indicators

Many of the data points we have discussed so far may not be very useful alone. Data that you collect must be aggregated and correlated with other types of data to create information. Data that is considered useful should also match measurements or metrics that have been previously developed by the organization. Metrics can be used to develop key indicators. Key indicators show overall progress toward goals or deficiencies that must be addressed. Key indicators come in four common forms:

- **Key performance indicators (KPIs)** Metrics that show how well a business process or even a system is doing with regard to its expected performance.
- **Key risk indicators (KRIs)** Can show upward or downward trends in singular or aggregated risk for a system, process, or other area of interest.
- **Key control indicators (KCIs)** Show how well a control is functioning and performing.
- **Key goal indicators (KGIs)** Overall indicators that may use the other indicators to show how well organizational goals are being met.

Most of these indicators are created by aggregating, correlating, and analyzing relevant security process data, to include both technical and administrative data. Examples of data that can be used to produce these metrics include vulnerabilities discovered during a technical scan, risk analysis, and summaries of user habits and incidents.

Management Review and Approval

An important data source that should be collected, and also kept for governance reasons, is the paper trail that results from management review and approval of assessments, reports, security actions, and so on. This information is important so that both the organization and any external entities understand the detailed process completed to approve a policy, implement a control, or respond to a risk or security issue. Not only is this documentation trail important for audits and governance, but it also establishes that the security program has management approval and involvement. It can also establish due diligence and due care in the event that any liabilities come from overall management of the security process.

REVIEW

Objective 6.3: Collect security process data (e.g., technical and administrative) In this objective we looked at a sampling of data points that a security program should collect—and the sources from which to collect them—to maintain and manage its security program. Data can come from a variety of sources, including technical data from applications, devices, log files, and SIEM systems. Relevant data can also come from written sources, such as

visitor access control logs, or in the form of process documentation such as infrastructure diagrams, configuration files, vulnerability assessment results, and so on. Even narrative data based on interviews from personnel can be useful. Types of data that are critical to collect, include, but are not limited to, account management data, backup verification data, and disaster recovery and response data. Metrics involve the use of specific data to create key indicators, which are specific data points of interest to management. Finally, a critical source of information comes from documentation reflecting management review and approval of security processes and actions.

6.3 QUESTIONS

1. You are an incident response analyst for your company. You are investigating an incident involving an employee who accessed restricted information he was not authorized to view. In addition to reviewing device and application logs, you also wish to establish the events for the timeline starting when the individual entered the facility and various other physical actions he took. Which of the following sources of information could help you establish these events in the incident timeline? (Choose two.)
   - **A.** Electronic badge access logs
   - **B.** Active Directory event logs
   - **C.** Closed-circuit television video files
   - **D.** Workstation audit logs
2. Which of the following is an aggregate of data points that can show management how well a control is performing or how risk is trending in the organization?
   - **A.** Quantitative analysis
   - **B.** Security process data
   - **C.** Metrics
   - **D.** Key indicators

6.3 ANSWERS

1. **A C** Badge access logs and video surveillance logs can help establish the events of the incident timeline that show the employee's physical activities.
2. **D** Key indicators are aggregates of data points that have been collected, correlated, and analyzed to represent important metrics regarding performance, goals, and risk in the organization.



Objective 6.4 Analyze test output and generate report

In this objective we will addresses what you should do with the test output and results, and how you should report those results. Different organizations and regulations have various reporting requirements, but we will discuss key pieces of the analysis and reporting process you should know for the CISSP exam in this objective, as well as remediation, exception handling, and ethical disclosure.

Test Results and Reporting

Once an organization has collected all relevant data from a test, it should analyze the test outputs and report those results to all stakeholders. Stakeholders include management, but may also include business partners, regulatory agencies, and so on. Each organization has its own internal reporting requirements, such as routing the report to different departments and managers, edits and changes that must be made, and reporting format. All of these requirements should be taken into consideration when reporting, but the key take-away here is that the primary stakeholders (i.e., senior management and other key organizational personnel) should be informed as to the results of the test, how it affects the security posture and risk of the organization, and what the path forward is for reducing the risk and strengthening the security posture.

Analyzing the Test Results

Analyzing the results of a test involves gathering all relevant data, as discussed in Objective 6.3, analyzing the results and output of tests that were executed during the evaluation, and distilling all this data into actionable information. Regardless of the evaluation methodology or the amount of data that goes into the analysis, properly evaluating the test results should provide the following:

- Analysis that is thorough, clear, and concise
- Reduction or elimination of repetitive or unnecessary data
- Prioritization of vulnerabilities or other issues according to severity, risk, cost, etc.
- Information that shows what is going on in the infrastructure and how it affects the organization's security posture and risk
- Root causes of any issues, when possible
- Business impact to the organization
- Mitigation alternatives or compensating controls
- Identifying changes to key performance indicators (KPIs), key goal indicators (KGIs), and key risk indicators (KRIs), as well as other metrics defined by the organization.

Cross-Reference

KPIs, KGIs, and KRIs were defined in Objective 6.3.

Reporting

After completing the test, security personnel report the results and any recommendations to management and other stakeholders. The goal of a report is to inform stakeholders of the actual situation with regard to any issues, shortcomings, or vulnerabilities that may have been uncovered during the test. The report should include historical analysis, root causes, and any negative trends that management should know about. Additionally, any positive results, such as reduced risk and good security practices, should also be highlighted in the report. The findings should be discussed in technical terms for those who have the knowledge and experience to implement mitigations for any discovered vulnerabilities, but often a nontechnical summary of the analysis is needed in the final report for senior management and other nontechnical stakeholders to understand.

The report should clearly convey the organization's security posture, compliance status, and risk incurred by systems or the organization, depending on the scope and context of the report. Relevant metrics that have been formally defined by the organization, such as the aforementioned indicator metrics, should also be reviewed. Finally, recommendations and other mitigations should be included in the report to justify expenditures of resources (money, time, equipment, and people) needed to mitigate any issues.

Remediation, Exception Handling, and Ethical Disclosure

In addition to analyzing test results and reporting, exam objective 6.4 requires that we examine three critical areas related to the results of evaluatons. In addition to other requirements detailed throughout this domain and other parts of the book, you must understand requirements for remediating vulnerabilities, handling exceptions to vulnerability mitigation management, and disclosing vulnerabilities as part of your ethical professional behavior.

Remediation

As a general rule, all vulnerabilities should be identified as soon as possible and mitigated in short order. This may mean patching software, swapping out a hardware component, requiring additional training for personnel, developing a new policy, or even altering a business process. It's generally not cost-effective for an organization to try to mitigate all discovered vulnerabilities at once; instead, the organization should prioritize the vulnerabilities according to several factors. Severity is a top priority, closely followed by cost to mitigate, level of actual risk to the organization and its systems, and scope of the vulnerability. For example, a vulnerability that presents a low risk to an organization because it only affects a single system that is not

connected to the outside world may be prioritized lower for remediation than a vulnerability that affects several critical systems and presents a higher risk.

Remediation actions should be carefully considered by management and included in the formal change and configuration management processes. Actions should also be formally documented in process or system maintenance records, as well as risk documentation.

Exception Handling

As mentioned earlier, vulnerabilities should be mitigated as soon as practically possible, based mainly on severity of the vulnerability. However, there are times when a vulnerability cannot be easily mitigated for various reasons, including lack of resources, system downtime, regulations, or other reasons. *Exception handling* refers to how vulnerabilities are handled when they cannot be immediately remediated. For example, discovering vulnerabilities in a medical device that cannot be easily patched due to U.S. Food and Drug Administration (FDA) regulations requires that the organization develop an exception-handling process to mitigate vulnerabilities by employing compensating controls.

The exception process should start with notifying the appropriate individuals who can make the decision regarding mitigation options, typically senior management; documenting the exception and the reasons for it; and determining compensating controls that can reduce the risk of not directly mitigating the vulnerability. There should also be a follow-up plan to look at the long-term viability of mitigating the vulnerability on a more permanent basis, which can include upgrading or replacing the system, changing the control, or even altering business processes.



EXAM TIP Understand how exceptions to policy, such as a vulnerability that cannot be immediately remediated, are handled through compensating controls, documentation, and follow-up on a periodic basis.

Ethical Disclosure

Ethical disclosure refers to a cybersecurity professional's ethical obligation to disclose the discovery of vulnerabilities to the organization's stakeholders. This ethical obligation applies whether you are an employee of the organization and discover a vulnerability during your routine duties or are an outside assessor employed to conduct an assessment or audit on an organization. In either case (or any other scenario), if you discover a vulnerability in a software or hardware product, you have an ethical obligation to disclose it. You should disclose any discovered vulnerabilities to organizations using the system or product, the creator/developers of the product, and, when necessary, the appropriate professional communities. As a professional courtesy, you should not disclose a newly discovered vulnerability to the general population without first disclosing it to those entities mentioned, since the vulnerability could be used by malicious entities to compromise systems before there is a mitigation for it.

REVIEW

Objective 6.4: Analyze test output and generate report This objective summarized what you should consider when analyzing and reporting evaluation results, to include general requirements for analyzing test output and reporting test results to all stakeholders. This objective also addressed vulnerability remediation, exception handling, and ethical disclosure of vulnerabilities.

6.4 QUESTIONS

1. You have discovered a vulnerability in a software product your organization uses. While researching patches or other mitigations for the vulnerability, you find that this vulnerability has never been documented. Which of the following should you do as a professional cybersecurity analyst? (Choose all that apply.)
   - **A.** Contact the software vendor directly to report the vulnerability.
   - **B.** Immediately post information about the vulnerability on public security sites.
   - **C.** Say nothing; if no one knows the vulnerability exists, then no one will attempt to exploit it.
   - **D.** Inform your supervisor.
2. After a routine vulnerability scan, you find that several critical servers have an operating system vulnerability that exposes the organization to a high risk of exploitation. The vulnerability is in a service that is not used by the servers but is running by default. There is a patch for the vulnerability, but it involves taking the servers down, which is not acceptable due to high data processing volumes during this time of the year. Which of the following is the best course of action to address this vulnerability?
   - **A.** Do nothing; since the servers don't use that particular service, the vulnerability can't affect them.
   - **B.** Take the servers down immediately and patch the vulnerability on each one.
   - **C.** Disable the service from running on the critical systems, and once the high processing times have passed, then patch the vulnerability.
   - **D.** Disable the service and do not worry about patching the vulnerability, since the servers don't use that service.

6.4 ANSWERS

1. **A D** You should contact the software vendor and report the vulnerability so that a patch or other mitigation can be developed for it. You should also contact your supervisor so that management is aware of the vulnerability and can determine any mitigations necessary to protect the organization's assets.

2. **C** The vulnerability must be patched eventually, but in the short term, simply disabling the service may mitigate or reduce the risk somewhat until the high data processing time has passed; then the servers can undergo the downtime required to patch the vulnerability. Doing nothing is not an option; even if the servers do not use that particular service, the vulnerability can be exploited. Taking down the servers immediately and patching the vulnerability on each one is generally not an option since it is a time of high-volume data processing and this may severely impact the business.



Objective 6.5 Conduct or facilitate security audits

To conclude our discussion on security assessments, tests, and audits, in this objective we will address conducting and facilitating security audits, specifically using internal, external, and third-party auditors. We will reiterate the definition of auditing as a process, as well as audits as distinct events, and provide some examples of audits and the circumstances under which different types of audit teams would be most beneficial.

Conducting Security Audits

In Objective 6.1, we briefly described audits under two different contexts. First, auditing is a detective control and a process that should take place on an ongoing basis. It is a security business process that involves continuously monitoring and reviewing data to detect violations of security policy, illegal activities, and deviations from standards in how systems and information are accessed. However, security audits can also be distinct events that are planned to review specific areas within an organization, such as a system, a particular technology, and even a set of business processes. A security audit can also be used to review events such as incidents, as well as to ensure that organizations are following standardized procedures in accordance with compliance or governance requirements.



NOTE For our purposes here, assume that "audits" and "auditors" refer throughout this objective as specifically "security audits" and "security auditors," unless otherwise specified.

For the purposes of this objective we will look at three different types of auditing teams that can be used to conduct an audit: internal auditors, external auditors, and third-party auditors. These are similar to the three types of assessment teams discussed in Objective 6.1, so you'll notice some overlap in this discussion.



EXAM TIP There is little difference between internal, external, and third-party assessors, also discussed briefly in Objective 6.1, and the three types of auditors discussed here, except where there are minor nuances in the types of assessments versus audits. The teams that can perform them have the same characteristics.

Internal Security Auditors

Internal auditors are likely the most cost-effective to use compared to external and third-party audit teams. Internal auditors work full time for the organization in question, so auditing is already part of their normal daily duties. Cybersecurity auditors are tasked with reviewing logs and other data on a daily basis to detect policy violations, illegal activities, and anomalies in the infrastructure. However, auditors can be assigned to specific events, such as auditing the account management process or auditing the results of an incident response or a business continuity exercise.

The advantages of using internal auditors include

- Cost effectiveness, since auditors already work for the organization
- Less difficulty in scheduling an internal audit versus an external audit
- Familiarity with the organization, its personnel, processes, and systems
- Reduces exposure of sensitive data to outsiders

There are, however, disadvantages to using internal audit teams, and indeed there are circumstances when this may not be permitted, particularly when the audit has been directed by an outside entity, such as a regulatory agency. Disadvantages of using an internal audit team include

- Influenced by organizational politics and conflicts of interest
- May not be full-time auditors and likely have other duties that must be performed during an audit event
- Lack of independence from organizational structure and management
- Not as experienced in a wide variety of auditing tools and techniques
- May not be allowed by external agencies if governance requires an independent audit

External Security Auditors

Just like an external assessment team, as mentioned in Objective 6.1, an external audit team normally works for a business partner entity of the organization. Contracts or other legal agreements between organizations may stipulate that certain systems, activities, or processes must be audited on a regular basis. An external auditor, also considered a second-party auditor, will periodically come into the organization and audit specific systems or processes to meet the requirements of the contract.

As with internal audit teams, using external auditors has disadvantages and advantages. Advantages include

- Cost may be more than internal audit team, but is typically already budgeted and built into the contract
- Not as easily influenced by organizational politics or conflicts of interest
- Defined schedule due to contract requirements (e.g., annually)

Disadvantages of using an external team include

- Lack of familiarity with the organization's people, internal processes, and infrastructure
- Lack of independence (allegiances to the business partner, not the organization)
- May be influenced by the business partner's internal conflicts of interest or politics
- May incur some of the same limitations as an internal team, such as split time between a regular IT or cybersecurity job, lack of access to advanced auditing tools and techniques, etc.

Third-Party Security Auditors

The third and final option relevant to this objective is the use of third-party auditors. Third-party auditors do not work for the organization or any of its business partners. A third-party audit team is normally required for independence from any internal stakeholder entity. This is the type of team that may be employed by a regulatory agency, for example, to ensure compliance with laws or regulations. Depending on the type of audit, an organization may not have a choice about whether to use a third-party audit team; it may be levied as part of external governance requirements.

As with the other two types of teams, using third-party audit teams has advantages and disadvantages. Advantages of using a third-party audit team include

- True independence from any organizational conflict of interest or politics
- Access to advanced auditing tools and techniques
- Establishes unbiased proof of compliance or noncompliance with regulations or standards

The disadvantages of using a third-party audit team include

- Most expensive option of the three types of audit teams; the expense cannot always be predicted or budgeted
- Sometimes difficult to schedule due to other auditing commitments, even if required on a recurring basis
- Lack of familiarity with the organization's personnel, processes, and infrastructure



EXAM TIP Remember that external auditors work for business partners or stakeholders outside of the organization. Third-party auditors work for independent organizations, such as regulatory agencies.

REVIEW

Objective 6.5: Conduct or facilitate security audits In this objective we discussed conducting security audits. Auditing is an ongoing business process that looks for wrongdoing and anomalies in operations. However, an audit is also an event used to review compliance standards for systems, processes, and other activities. Audits can be performed by one of three types of audit teams: internal teams, external teams, and third-party teams. Internal teams are more cost-effective but lack independence from the organization and may not have access to the right audit tools and techniques. External teams work for a business partner or other stakeholder and audit processes and activities as required by a contract. Third-party audit teams are more expensive but allow some level of independence from organizational stakeholders and may be required in the event of regulatory audits.

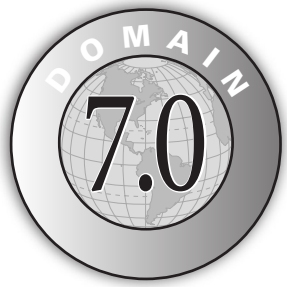
6.5 QUESTIONS

1. You are a cybersecurity analyst for your company. You have been tasked with auditing account management in another division of the company. Which of the following types of auditing would this be considered?
   - **A.** Third-party audit
   - **B.** External audit
   - **C.** Internal audit
   - **D.** Second-party audit
2. Your company must be audited for compliance with regulations that protect healthcare information. Which of the following would be the most appropriate type of auditors to perform this task?
   - **A.** External auditors
   - **B.** Internal auditors
   - **C.** Second-party auditors
   - **D.** Third-party auditors

6.5 ANSWERS

1. **C** An internal audit uses auditors from within an organization to assess another part of the organization.
2. **D** For a compliance audit, third-party auditors are usually the most appropriate type of audit team to conduct the effort, due to their independence.

Security Operations



Domain Objectives

- **7.1** Understand and comply with investigations.
- **7.2** Conduct logging and monitoring activities.
- **7.3** Perform Configuration Management (CM) (e.g., provisioning, baselining, automation).
- **7.4** Apply foundational security operations concepts.
- **7.5** Apply resource protection.
- **7.6** Conduct incident management.
- **7.7** Operate and maintain detective and preventative measures.
- **7.8** Implement and support patch and vulnerability management.
- **7.9** Understand and participate in change management processes.
- **7.10** Implement recovery strategies.
- **7.11** Implement Disaster Recovery (DR) processes.
- **7.12** Test Disaster Recovery Plans (DRP).
- **7.13** Participate in Business Continuity (BC) planning and exercises.
- **7.14** Implement and manage physical security.
- **7.15** Address personnel safety and security concerns.

Domain 7 is unique in that it has the most objectives of any of the CISSP domains, and it accounts for approximately 13 percent of the exam questions. You'll find that many of the objectives covered in this domain, Security Operations, have also been briefly discussed throughout the entire book. This is because security operations are diverse and overarching activities that span multiple areas within security.

In this domain we will examine a wide range of subjects, including those that are reactive in nature, such as investigations, logging and monitoring, vulnerability management, and incident management. We will also look at the details of how to ensure a secure change and configuration management process that is supported by patch management.

We will review some of the foundational security operations concepts that we also discussed in Domain 1 and apply some of those concepts to resource protection. We will also look at some of the more technical details of detective and preventive measures, such as firewalls and intrusion detection systems. Four of the objectives address business continuity planning and disaster recovery, and we will discuss the strategies involved with each topic as well as how to implement and test the plans associated with these processes. Finally, we will review physical and personnel safety and security concerns.



Objective 7.1 Understand and comply with investigations

In Objective 1.6 we briefly touched on the types of investigations you may encounter in security, and we also reviewed the related topics of legal and regulatory issues in Objective 1.5. These two objectives go hand-in-hand with Objective 7.1, which carries our discussion a bit further by focusing on *how* investigations are conducted.

Investigations

Recall from Objective 1.6 that the four primary types of investigations are administrative, regulatory, civil, and criminal investigations. Regardless of the type of investigation, however, most of the activities, processes, and techniques that are used are common across all of them. This includes how to collect and handle evidence; reporting and documenting the investigation; the investigative techniques that are used; the digital forensics tools, tactics, and procedures that are implemented; and the artifacts that are discovered on computing devices using those tools, tactics, and procedures. These common activities are the focus of this objective.

Cross-Reference

The types of investigations you may encounter in cybersecurity were discussed at length in Objective 1.6.

Forensic Investigations

Computers, mobile devices, network devices, applications, and data files all contain potential evidence. In the event of an incident involving any of them, they must all be investigated. Computing devices can be part of an incident in three different ways:

- As the target of the incident (e.g., as an attack on a system)
- As the tool of the incident (used to directly perpetrate a crime, such as a hacking attack)
- Incidental to the event (part of the event but not the direct target or tool used, such as researching how to commit a murder, for instance)

Computer forensics is the science of the identification, preservation, acquisition, analysis, discovery, documentation, and presentation of evidence to a court of law (either civil or criminal), corporate management, or even to a routine customer. Computer forensic investigations are carried out by personnel who are uniquely qualified to perform them. *E-discovery* is another term with which you should be familiar; it is the process of discovering digital evidence and preparing that information to be presented in court. The most important part of the computer forensics process is how evidence is collected and handled, discussed next.

Evidence Collection and Handling

Evidence preservation is the most crucial part of an investigation. Even if an unskilled investigator is not properly trained to analyze evidence, the investigation can be saved by ensuring that the evidence is preserved and protected at all stages of the investigation. Evidence preservation involves proper collection and handling, including chain of custody, physical protection, and logical protection. We will discuss proper collection and handling procedures in the next few sections. However, chain of custody is an important one to discuss first.

Chain of custody refers to the consistent, formalized documentation of who is in possession of evidence at all stages of the investigation, up to and including when the evidence is presented in court or another proceeding. Chain of custody starts at the time evidence is obtained. The individual collecting the evidence initiates a chain of custody form that describes the evidence and documents the time and location it was collected, who collected it, and its transfer. From then on, every individual who takes possession of the evidence, relocates it, or pulls it out of storage for analysis must document those actions on the chain of custody form. This ensures that an uninterrupted sequence of events, timelines, and locations for the evidence is maintained throughout the entire investigation. Chain of custody assures evidence integrity and protects against tampering. This is one of the most crucial pieces of documentation you must initiate and maintain throughout the entire investigative cycle.

Evidence Life Cycle

Evidence has a defined life cycle. This means that the moment evidence is identified during the initial response to the investigation, it is collected and handled carefully, according to strict procedures. Chain of custody begins this critical process, but it doesn't end with that event. Evidence must be protected throughout the investigation against tampering; even the

appearance of intentional or inadvertent changes to evidence may call its investigative value and admissibility into court into question.

The evidence life cycle consists of four major phases, summarized as follows:

- **Initial response** Evidence is identified and protected; a chain of custody is initiated.
- **Collection** Evidence is acquired through forensically sound processes; evidence integrity is preserved.
- **Analysis** All evidence is analyzed at the technical level to determine the timeline and chain of events. Determining innocence or guilt of a suspect is the goal as well as identifying the root cause of the incident.
- **Presentation** Evidence is summarized in the correct format for presentation and reporting for corporate management, a customer, or to a court of law.



NOTE This life cycle, as with all other life cycles, may be different depending upon the methodology or standard used, since there are many different life cycles that exist in the investigative world. However, all of them agree on fundamental evidence collection and handling processes, which are standardized all over the world.

Obviously, there are more in-depth processes and procedures that must take place at each of these phases, and we will discuss those at length in the next section. Figure 7.1-1 summarizes the evidence life cycle.

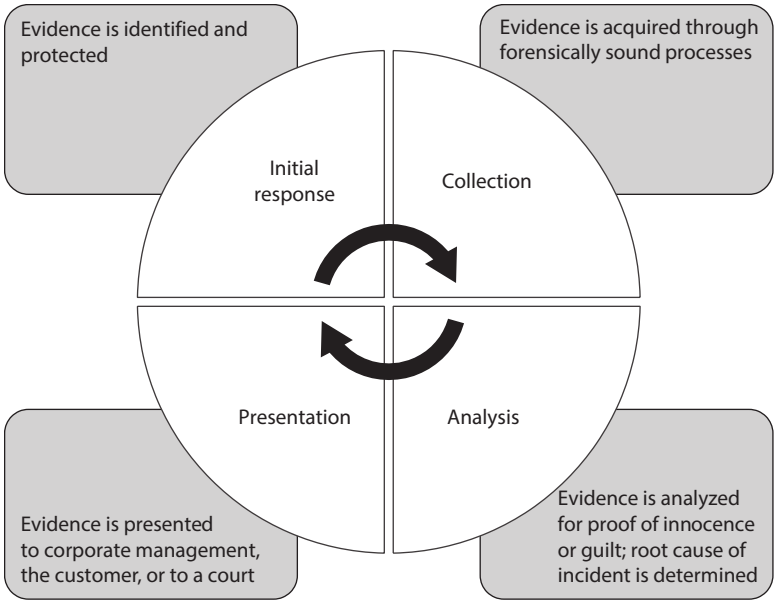


FIGURE 7.1-1 The evidence life cycle

Evidence Collection and Handling Procedures

There are standardized evidence collection and handling procedures used throughout the world, regardless of the type of investigation. These have all been adopted as formal standards by different law enforcement agencies, security firms, and professional organizations. They are summarized as follows:

- Secure the scene of the crime or incident against unauthorized personnel.
- Photograph the scene before it is disturbed in any way.
- Don't arbitrarily power off devices until any available live evidence has been gathered from them.
- Obtain legal authorization from law enforcement or an asset owner before removing items from the scene.
- Inventory all items removed from the scene.
- Transport and store all evidence items in protected containers and store them in secure areas.
- Maintain a strong chain of custody at all times.
- Don't perform a forensic analysis on original evidence items; forensically duplicate the evidence item and perform an analysis on the duplicate to avoid destroying or compromising the original.



EXAM TIP Once evidence is obtained from the source, such as a device, logs, and so on, that source may be placed on what is known as *legal hold*. Legal hold ensures that any devices or media that contain the original evidence must be kept in secure storage, and access must be controlled. These items cannot be reused, destroyed, or released to anyone outside the chain of custody until cleared by a legal department or court.

Artifacts

Artifacts are any items of potential evidentiary value obtained from a system. They are usually discrete pieces of information in the form of files, such as documents, pictures, executables, e-mails, text messages, and so on that are found on computers, mobile devices, or networks. However, they can also be information such as screenshots, the contents of RAM, and storage media images.

Artifacts are used as evidence of activities in investigations and can serve to support audit trails. For example, the Internet history files from a computer can support an audit log that indicates an individual visited a prohibited website. Files such as pictures or documents can indicate whether individuals are performing illegal activities on their system.

Note that artifacts by themselves are not indicative of an individual's guilt or innocence; the presence of artifacts on a system mutually corroborates audit trails and other sources of information during an investigation. Artifacts must be investigated on their own merit before they are determined to meet the requirements of evidence. As discussed earlier in the objective, digital artifacts, as potential evidence, must be collected and handled with care.

Digital Forensics Tools, Tactics, and Procedures

While knowledge and experience with evidence collection and handling procedures are among the most crucial skills forensic investigators can have, they should also have core technical analysis skills and knowledge of a variety of subjects, including how storage media is constructed and operates, operating system architecture, networking, programming or scripting, and security. If you are conducting a forensic investigation, these skill sets will assist you greatly when performing some of the following forensic tasks:

- Data acquisition from volatile memory or hard drives using forensic techniques
- Establishing and maintaining evidence integrity through hashing tools to ensure artifacts are not intentionally or inadvertently changed
- Data carving (the process used to "carve" discrete data or files from raw data on a system using forensics processes) to locate and preserve artifacts that have been deleted or hidden

We are long past the days when computer forensics was performed mostly on simply end-user desktop computers or servers. In today's environment, devices and data are integrated all the way from end-user mobile devices, into the cloud, and back to the organization's infrastructure. While core forensics knowledge and skills are still necessary, so too are knowledge and skills related to specific procedures that are tied to more narrowly focused areas within forensics. These areas often require specialized knowledge and tools in addition to generalized forensics skills. These areas of expertise include

- Cloud forensics
- Mobile device forensics
- Virtual machine forensics

The choice of tools that a forensic investigator uses is important. There are specific tools that are used for specific actions, including data acquisition, log aggregation review, and so on. Each investigator or organization typically has favorite tools they use to perform all of these tasks. Some tools are proprietary, commercial-off-the-shelf enterprise-level software suites sold specifically for forensics processes, but many are simply individual tools that come with the operating system itself, such as utilities or built-in applications. Many forensics tools may also be internally developed utilities, to include scripts, for instance, or even open-source software utilities or applications.

Regardless of which digital forensics tools an organization uses, the following are some key things to remember about a forensics tool set:

- The tools should be standardized and thoroughly documented. An organization should have established procedures for the investigator to follow when using a tool.
- Forensics tools are often validated by professional organizations or national standards agencies; these tools should be preferred over tools developed in-house or tools whose origins and effectiveness can't be easily verified.
- Tools that can offer repeatable and verifiable results should be used; if a tool does not image the same hard drive consistently every single time, for example, then its usefulness may be limited since its integrity cannot be trusted.

While this objective can't possibly cover every single tool available to you during your forensic investigation, you can generally categorize tools in the following areas:

- Network tools (protocol analyzers or sniffers such as Wireshark and tcpdump)
- System tools (used to obtain technical configuration information for a system)
- File analysis tools
- Storage media imaging tools (both hardware and software tools)
- Log aggregation and analysis tools
- Memory acquisition and analysis tools
- Mobile device forensics tools

Investigative Techniques

Investigations attempt to discover what happened before, during, and after an incident. The goal of investigators is to identify the root cause of an incident and help ensure someone is held accountable for illegal acts or those that violate policy. Investigators also want to answer questions such as the who, what, where, when, and how of an incident. Investigators have the primary tasks of

- Collecting and preserving all evidence
- Determining the timeline and sequence of events
- Determining the root cause of and methods used during an incident
- Performing a technical analysis of evidence
- Submitting a complete, comprehensive, unbiased report

Investigators should always treat all investigations as if the results will eventually be presented in a court of law. This is because many investigations, even ones that seemingly start out as innocuous policy violations, may go to court if the evidence indicates that a violation of the law has occurred.

During an investigation, there are key points to keep in mind:

- Remain unbiased; don't go into an investigation automatically presuming guilt or innocence.
- Always have another investigator validate work you have performed.
- Maintain documented, verifiable forensic procedures.
- Keep all investigation information confidential to the extent possible; it should only be shared with key personnel such as senior corporate managers or law enforcement.
- Only perform procedures on evidence that you are trained and qualified to perform; don't undertake any actions for which you are not qualified.
- Ensure you have the proper tools to perform forensic activities on evidence; using the wrong tool could inadvertently destroy or compromise the integrity of evidence.

Reporting and Documentation

Documentation is one of the most critical aspects of investigations. An investigator should document every action they take during an investigation. The documentation for investigations should meet legal requirements and be thorough and complete. The content of documentation includes

- All actions involving evidence and witnesses (chain of custody, artifacts collected, witness interviews, etc.)
- Dates, times, and relevant events
- All forensic analysis of evidence

Reports and relevant documentation are usually delivered formally to the corporate legal department, human resources, or lawyers for all parties, as well as law enforcement investigators. Reports and investigation documentation should be clear and concise and present only the facts regarding an incident.

A good report includes investigative events, timelines, and evidence. The analysis portion of the report includes determination of the root cause(s), attack methods used during the incident, and the assertion of proof of guilt or innocence of the accused.

Investigative reports usually are formatted according to the desires of the corporate management or the court or agency that maintains jurisdiction over the investigation. In general, however, the investigation report should consist of an executive summary, the details of the events of the investigation, any findings and supporting evidence, and conclusions regarding the root cause of the case. Characteristics of a well-written investigative report include

- Clear, concise, and nontechnical
- Well written and well researched
- Answers to critical investigative questions of who, what, where, when, why, and how

- Conclusions supported by evidence
- Unbiased analysis

In addition to documentation and reporting, investigations also may make use of witness depositions or testimony. Witnesses are often asked to testify if they have direct knowledge of the facts of the case. Investigators can also be required to testify in court to detail the facts of the investigation.

REVIEW

Objective 7.1: Understand and comply with investigations This objective provided an opportunity to discuss details of investigations in more depth. Whereas Objective 1.6 covered the different types of investigations, in this objective we discussed the details of *how* investigations are conducted. We examined in particular forensic investigations, which involve gathering evidentiary data from computing systems.

Evidence collection and handling is the most crucial part of an investigation, since once evidence has been destroyed or compromised, it may not be recovered or trusted. The evidence life cycle consists of four general phases: initial response, collection, analysis, and presentation. The most critical part of evidence collection and handling is to establish a chain of custody that follows the evidence over its entire life cycle. Chain of custody assures that the evidence is always accounted for during transfer, storage, and analysis, and helps to rebut claims that the evidence has been tampered with or is unreliable. Other critical evidence handling activities include securing the scene of the incident or crime, photographing all evidence before it is removed, securely transporting and storing the evidence, and conducting analysis using only verifiable forensic procedures.

Artifacts are any items of potential evidentiary value obtained from a system, including files, logs, screen images, media, network traffic, and the contents of volatile memory. Artifacts are used to support a legal case and corroborate with other sources of information.

Digital forensics consists of a wide variety of tools, techniques, and procedures. The forensic investigator should be well-versed in a variety of disciplines, including networking, operating systems, programming, and other specific areas such as cloud computing, mobile device forensics, and virtual machine technology. Digital forensics tools can be categorized in terms of network tools, system tools, file analysis tools, storage media imaging tools, log aggregation analysis tools, memory acquisition and analysis tools, and mobile device tools.

Investigative techniques include solid knowledge of legal and forensic procedures with regard to evidence collection and handling, as well as technical areas of expertise. Investigators should also understand how to present an analysis of evidence in a court of law and should conduct all investigations as if they will proceed in that direction. Investigators should also approach every incident with an open mind with no bias as to the guilt or innocence of a suspect.

Forensic reports and documentation must be thorough and complete; they must follow the format prescribed by the corporate entity, customer, or the court of jurisdiction. They should include an executive summary, technical findings, and analysis of the evidence that supports those findings. They should also propose a conclusion and list any relevant facts pertinent to the case. Reports should also be clear and understandable to nontechnical personnel.

7.1 QUESTIONS

1. You have been called to investigate an incident of an employee who has violated corporate security policies by downloading copyrighted materials from the Internet. You must collect all evidence relating to the incident for the investigation, including the employee's workstation. Which one of the following is the most critical aspect of the response?
   - **A.** Establishing a chain of custody
   - **B.** Analyzing the workstation's hard drive
   - **C.** Creating a forensic duplicate of the workstation's hard drive
   - **D.** Creating a formal report for management
2. Which of the following best describes one of the primary tasks a forensic investigator must complete?
   - **A.** Ensuring that the evidence proves a suspect is guilty
   - **B.** Determining a timeline and sequence of events
   - **C.** Performing the investigation alone to ensure confidentiality
   - **D.** Manually analyzing device logs

7.1 ANSWERS

1. **A** During the initial response, creating a solid chain of custody is critical for evidence integrity and preservation. The other choices refer to processes that normally take place after the initial response.
2. **B** One of the primary tasks of the investigator is to determine a timeline and sequence of events that occurred during an incident. The other choices indicate things an investigator should not do, such as only looking for evidence that proves guilt, performing an investigation alone, or manually analyzing logs.



Objective 7.2 Conduct logging and monitoring activities

This objective covers the more technical aspects of logging and monitoring the network infrastructure and traffic. Although we have touched on these topics throughout the book, this objective addresses the need for and the process of collecting data from different sources all over the network, aggregating that data, and then performing analysis and correlation to determine the overall security picture for the network.

Logging and Monitoring

Remember that in Objective 6.1 you learned about security audits; auditing is directly enabled by logging network activity and monitoring the infrastructure for negative events and anomalous behavior. However, there's more to these critical activities than that. Logging and monitoring are necessary to maintain understanding and visibility of what's going on in the network infrastructure, which includes network devices, traffic, hosts, their applications, and user activity. You need to understand not only what's going on in the network at any given moment, but also what's happening over time, so that you can perform historical analysis and predict potentially negative trends.

Monitoring includes not only security monitoring but also performance monitoring, function monitoring, and user behavior monitoring. The purpose for all of this logging and monitoring is to collect small pieces of data that, when put together and given context, generate information that enables you to take proactive measures to defend the network. Some key elements of the infrastructure that you must monitor include

- Network devices and their performance
- Servers
- Endpoint security
- Bandwidth utilization and network traffic
- User behavior
- Infrastructure changes or departures from normal baselines

Much of this information is generated by logs, particularly from network and host devices. Logs that cybersecurity analysts review on almost a daily basis include firewall logs, proxy logs, and intrusion detection and prevention system logs. In this objective we will discuss many of the technologies that enable logging and monitoring, as well as how they are implemented.

Cross-Reference

Objective 7.7 provides a broader overview of firewalls and intrusion detection and prevention systems.

Continuous Monitoring

Continuous monitoring requires a resilient infrastructure that is able to collect, adjust, and analyze data on multiple levels, including both network-based data (e.g., traffic characteristics and patterns) and host-based data (such as host communications, processes, applications, and user activity). Continuous monitoring involves the use of IDS/IPSs and security information and event management systems (SIEMs).

We discuss continuous monitoring here in two different contexts. The first is more relevant to logging and monitoring the infrastructure and involves proactive monitoring of both the network infrastructure and its connected hosts to detect anomalies in configured baselines, as well as potentially malicious activities. The second context is not as technical but equally as important: monitoring overall system and organizational risk. Risk is monitored and measured on a continual basis so that any changes in the organization's risk posture can be quickly identified and adjusted if needed. Risk changes frequently due to several factors, which include the threat landscape, the organization's operating environment, technologies, the industry or market segment, and even the organization itself. All of these risk factors must be monitored to ensure risk does not exceed appetite or tolerance levels for the organization.

Intrusion Detection and Prevention

Historically, network and security personnel focused on simple intrusion detection capabilities. Modern security devices function as both intrusion detection and prevention systems (IDS/IPSs). They not only can detect and categorize a potential attack, but can also take actions to halt traffic that may be malicious in nature. An IDS/IPS may perform this function by dynamically rerouting network traffic, shunting connections, or even isolating hosts.

Intrusion detection typically relies on a combination of one or more of three models to detect problems in the infrastructure:

- **Signature- or pattern-based (rule-based) detection** Uses well-known signatures or patterns of behavior that match attack characteristics (e.g., traffic inbound on a particular port from a specific domain or IP address)
- **Anomaly- or behavior-based analysis** Detects deviations in normal behavior baselines
- **Heuristic analysis** Also detects deviations in normal behavior baselines, but matches those deviations to potential attack characteristics



EXAM TIP There is a difference between behavior-based analysis and heuristic analysis, although they are very similar. With behavior-based analysis, any deviations in the normal behavior baseline will be flagged and security personnel will be alerted. However, even deviations can be explained under certain circumstances. Heuristic analysis takes it a step further by looking to see what those abnormal behaviors might do, such as accessing protected memory areas, changing operating system files, or writing data to a hard drive.

Security Information and Event Management

With properly configured monitoring and logging, a large network, or even a medium-sized network, likely collects millions of pieces of data daily from a variety of network sources. These pieces of data come from hosts, network devices, user activity, network traffic, and so on. It would be impossible for a single person or even several people to sift through the data to make sense out of it and gain meaningful information from it. A good amount of data that comes from logging and monitoring may be insignificant; it's up to a security analyst to determine which of the millions of pieces of data are important and what they mean to the overall security of the network.

Fortunately, the daunting process of collecting, aggregating, correlating, and analyzing this data is not simply left up to human beings to perform. This is where automation significantly contributes to the security process. We have already discussed how security tool automation is critical to the security process, but here we are talking specifically about security information and event management (SIEM) systems.

A SIEM system is a multifunctional security device whose purpose is to collect data from various sources, aggregate it, and assist security analysts in analyzing it to produce actionable information regarding what is happening on the network. SIEM systems are often the central data collection point for all log files, traffic captures, and other forms of data, sometimes disparate. A SIEM system ingests all of this data and correlates seemingly unrelated data to connect data points and show how they actually do relate to each other. This device helps you to make intelligent, risk-based decisions in almost real time about the security of the network.

SIEM systems use a concept known as a *dashboard* to display information to security analysts and allow them to run extensive queries on information to get very detailed analysis from all of these different sources of information.

Egress Monitoring

Egress monitoring specifically examines traffic that is leaving the network. Egress monitoring is typically performed by firewall, proxy, intrusion detection, or data loss prevention systems. For the most part this will be routine traffic, but egress monitoring looks for specific security issues. Obviously, a major issue is malware. Often an attack may come in the form of a distributed denial-of-service (DDoS) attack carried out by a botnet that uses the network against itself by infecting different hosts, which then attack other hosts on the network or even hosts on an external network. Egress monitoring looks for signs that internal hosts have been compromised and are being controlled by an external malicious entity and are communicating with it.

In addition to malware, another issue egress monitoring is useful for detecting is data exfiltration. This usually involves sensitive data that is being illegally sent outside the network, in an uncontrolled manner, to unauthorized entities. Egress monitoring uses several different technologies to detect this issue; in addition to data loss prevention (DLP) technologies deployed on both network devices and user endpoints, security devices such as firewalls implement rule sets that look for large volumes of data as well as files with particular extensions, sizes, and other characteristics.

Cross-Reference

Data loss prevention was discussed in Objective 2.6.

Log Management

If an organization does not monitor its logs and react to them properly, the logs serve no useful function. Given that there may be thousands of devices writing logs, managing logs can seem like a daunting task. Again, this is where automation comes into play. Logs are usually automatically sent to central collection points, such as the aforementioned SIEM system, or even a *syslog* server, for examination. Often, manual log review must occur to solve a particular problem, research a specific event, or gain more details about what is going on with the network. However, these are usually the exceptions, and most of the log management process can be automated, as mentioned previously.

Most devices generate what are known generically as *event logs*. An event log records an occurrence of an activity or happening. An event is usually something that is considered on a singular basis and has definable characteristics. Basic information for an event in a log includes

- Event definition
- System or resource the event effects
- Identifying information for a host, such as hostname, IP address, or MAC address
- The user or other entity that initiated or caused the event
- Date, time, and duration of the event
- Event action (e.g., file deletion, privilege use, etc.)



EXAM TIP You should be familiar with the general contents of an event log entry, which typically includes an event definition, the system affected, host information, user information, the action that was taken, and the date and time of the event.

Log analysis, also primarily an automated task performed by SIEM systems, has the goal of looking through various logs to connect data points and ascertain any patterns between those aggregated data points.

Threat Intelligence

Threat intelligence is the process of collecting and analyzing raw threat data in an effort to identify potential or actual threats to the organization. This may involve determining threat trends to predict what a threat will do, historical analysis of threat data to recognize what happened during a particular event, or behavioral analysis to understand how a threat reacted under certain circumstances to the environment.

Note that the terms *threat data* and *threat intelligence* are similar but not the same thing. Threat data refers to raw pieces of information, typically without context, which may or may not be related to each other. An example is an IP address or a log entry that shows a connection between two hosts. Threat data only becomes threat intelligence when it is analyzed and correlated to gain useful insight into how the data relates to the organization's assets. Threat intelligence can come from various sources, called *threat feeds*, which include open-source, proprietary, and closed-source information.

Characteristics of Threat Intelligence

The effectiveness of threat intelligence can be evaluated based on the following three different characteristics, which will determine the quality or usefulness of the intelligence to the organization:

- **Timeliness** The intelligence must be obtained as soon as it is needed to be of any value in countering a threat.
- **Accuracy** The intelligence must be factually correct, accurate, and not contribute to false negatives or positives.
- **Relevance** The intelligence must be related to the given threat problem at hand and considered in the correct context when viewed with other factors.

Two other characteristics of threat intelligence are

- **Threat rating** Indicates the threat's potential danger level. Typically, the higher the rating, the more dangerous the threat.
- **Confidence level** The trust placed in the source of the threat intelligence and the belief that the threat rating is accurate.

Both threat ratings and confidence levels can be expressed on qualitative scales, from least dangerous to most dangerous, for example, or least level of confidence to highest level of confidence, respectively. Often threat ratings and threat confidence levels directly relate to the sources from which we gain intelligence, as some are more dependable than others. Threat intelligence sources are discussed next.



EXAM TIP Make sure you are familiar with the characteristics of threat intelligence timeliness, accuracy, and relevance, and that you understand the concepts of threat rating and confidence level.

Open-Source Intelligence

Open-source intelligence (OSINT) comes from sources that are available to the general public. Examples include public databases, websites, and general news. While open-source intelligence is very useful, it is typically broader and describes very general characteristics of threats,

which may not apply to your particular assets, vulnerabilities, or overall organization. Open-source intelligence comes in great volumes, which must be reduced, sorted, prioritized, and analyzed to determine its relevance to the organization. (Threat modeling, discussed a bit later, is useful for distilling OSINT.)

Closed-Source Intelligence

Closed-source intelligence comes from threat feeds that may be restricted in their availability. Consider classified government intelligence feeds, for example. These are not readily available to the general public due to data sensitivity or the sensitivity of their source, such as from an agent operating covertly in a foreign country or obtained with secret technology. Another key differentiator for closed-source intelligence versus OSINT is that typically closed-source intelligence is more accurate, more thoroughly authenticated, and holds a higher confidence level. Closed-source intelligence also often provides greater detail and fidelity about the threat, particularly as the intel is often focused on specific organizations, assets, and vulnerabilities that are targeted.

Proprietary Intelligence

Proprietary intelligence can be thought of as a closed-source intelligence feed, but it is usually developed by a private organization and sold, via subscription, to any organization that wishes to purchase it. This makes it more of an intelligence commodity as opposed to being restricted from the general public based on sensitivity. Many organizations purchase proprietary threat intelligence feeds from other companies, sometimes tailored to their specific market or circumstances.

Threat Hunting

Threat hunting is the active effort to determine whether various threats exist in an infrastructure. In some cases, an analyst may be looking to determine if specific threats or threat actors have already infiltrated the infrastructure and continue to maintain a presence. In other cases, threat hunting is more geared toward looking for a variety of threats on a continual basis to ensure that they don't ever get into the infrastructure in the first place. Threat hunting uses both threat intelligence feeds and threat modeling to determine more precisely which threats are more likely to target which assets in the infrastructure, rather than looking for generic threats. Then the threat hunters make a concerted effort to look for those specific threats or threat actors in the network.

Threat Modeling Methodologies

Several formalized methodologies have been developed to address the different characteristics and components of threats. Some address threat indicators, some address attack methods that threat sources can use against organizations (called *threat vectors*), and some allow for in-depth threat modeling and analysis. All these methodologies allow the organization to formally manage threats and are critical components of the threat modeling process. A few examples are listed and described in Table 7.2-1.

TABLE 7.2-1 Various Threat Modeling Methodologies

Threat Model	Description
Structured Threat Information eXpression (STIX)	Describes threat intelligence in a common language to facilitate its exchange between organizations
Trusted Automated eXchange of Indicator Information (TAXII)	Cyberthreat standard that describes how threat data can be shared; uses standard communications APIs to make it compatible with a variety of cyberthreat systems
OpenIOC	An open-source cyberthreat-sharing framework used to share threat data with other entities using XML
MITRE ATT&CK Framework	Public knowledge database of threat tactics and techniques
Diamond Model of Intrusion Analysis	Analytical model used to view the characteristics of threat actors/events and assists in the analysis to defend against them
Cyber Kill Chain	Cybersecurity model originally developed by Lockheed Martin to identify the various stages of threats during a cyberattack
Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE)	Methodology developed by Carnegie Mellon University's Software Engineering Institute (SIE) that focuses on operational risk, security controls, and security technologies in an organization
Trike	Open-source threat modeling methodology focused on auditing
Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege (STRIDE)	Threat modeling methodology created by Microsoft for incorporating security into application development by categorizing each threat into one of six categories: spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege.
Visual, Agile, and Simple Threat modeling (VAST)	Threat modeling methodology incorporated into the software development life cycle (SDLC) and frequently used in Agile development models
Process for Attack Simulation and Threat Analysis (PASTA)	Threat modeling methodology focused on integrating technical requirements with business process objectives

While an in-depth discussion on any of these threat methodologies is beyond the scope of this book, you should have basic knowledge about them for the CISSP exam.

Cross-Reference

We also discussed threat modeling in Objective 1.11.

User and Entity Behavior Analytics

User and entity behavior analytics (UEBA) focuses on patterns of behavior from users and other entities (e.g., service accounts or processes). UEBA goes beyond simply reviewing a log of user actions; it looks at user behavioral patterns over time to detect when those patterns of

behavior change. These behavioral patterns include when a user normally logs on or off of a system, which resources they access, and how they interact with the system as a whole.

When a pattern of behavior deviates from the normal baseline, it may be an indicator of compromise (IoC). It may indicate one of several possibilities that merit further investigation, such as:

- The user is violating a policy or doing something illegal
- The system itself is functioning but performing in a less than optimal manner
- The system is under attack from a malicious entity

As with all the other types of data in the infrastructure, user behavior data must be initially collected, aggregated, and analyzed to determine normal baselines of behavior.

REVIEW

Objective 7.2: Conduct logging and monitoring activities In this objective we discussed details of technical logging and monitoring. Logging and monitoring contribute to the auditing function by providing data to connect events to entities. Logs can come from various sources, including network devices, hosts, applications, and so on. Event log entries normally include details regarding the user that initiated the event, the identifying host information, a description or definition of the event, the date and time of the event, and what actually happened.

Continuous monitoring is a proactive way of ensuring that you not only have continuous visibility into what is happening in the network but also are able to perform historical analysis and trend prediction. Continuous monitoring also means the organization is continually monitoring its risk posture.

Intrusion detection and prevention systems use three methods of detection, sometimes in combination with each other: signature or pattern-based detection, behavioral-based detection, and heuristic detection.

Logs and other data from across the infrastructure can be fed into an automated system that aggregates and correlates all of this information, known as a SIEM system. SIEM systems allow instant visibility into the security posture of the network through dashboards and complex queries.

Egress monitoring allows security personnel to detect malware attacks that may make use of botnets and cause hosts to attack each other or, worse, attack external networks not owned by the organization. Egress monitoring also allows the organization to detect data exfiltration through secure device rule sets and data loss prevention systems.

Log management means that administrators actually review logs to detect malicious events, poor network performance, or negative trends. Most modern log management is automated through SIEM systems.

We also revisited and expanded upon the basic concepts of threat modeling. Threat modeling goes beyond simply listing generic threats that could be applicable to any organization; threat modeling takes a more in-depth, detailed look at how specific threats may affect an organization's assets and vulnerabilities. Threat modeling uses threat intelligence that is timely, relevant, and accurate, and that intelligence may come from a variety of threat feeds, such as open-source, closed-source, or proprietary sources. Various threat management and modeling methodologies exist, including STRIDE, VAST, PASTA, and many others.

Finally, we examined user and entity behavior analytics (UEBA), which looks for abnormal behavioral patterns from users, system accounts, and processes. These deviations of normal behavior patterns could indicate an issue with a user, the system, or an attack.

7.2 QUESTIONS

1. You are designing a new intrusion detection and prevention system for your company. You want to ensure that it has the capability to accept security feeds from the system's vendor to allow you to detect intrusions based on known attack patterns. Which one of the following detection models must you include in the system design?
   - **A.** Behavior-based detection
   - **B.** Heuristic detection
   - **C.** Signature-based detection
   - **D.** Intelligence-based detection
2. You are a cybersecurity analyst who works at a major research facility. As part of the organization's effort to perform threat modeling for its systems, you need to look at various proprietary intelligence feeds and determine which ones would be most likely to help in this effort. Which of the following is not an important characteristic of threat intelligence you should consider when selecting threat feeds?
   - **A.** Timeliness
   - **B.** Methodology
   - **C.** Accuracy
   - **D.** Relevance
3. Nichole is a cybersecurity analyst who works for O'Brien Enterprises, a small cybersecurity firm. She is recommending various threat methodologies to one of her customers, who wants to develop customized applications for Microsoft Windows. Her customer would like to incorporate a threat modeling methodology to help them with secure code development. Which of the following should Nichole recommend to her customer?
   - **A.** PASTA
   - **B.** TRIKE
   - **C.** VAST
   - **D.** STRIDE

7.2 ANSWERS

1. **C** Signature-based detection allows the system to detect attacks based on known patterns or signatures.
2. **B** Methodology is not a consideration in evaluating intelligence feeds. To be useful to the organization, threat intelligence should be timely, relevant, and accurate.
3. **D** STRIDE (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege) is a threat modeling methodology created by Microsoft for incorporating security into application development. None of the other methodologies listed are specific to application development, except for VAST (Visual, Agile, and Simple Threat Modeling), but it is not specific to Windows application development.



Objective 7.3 Perform Configuration Management (CM) (e.g., provisioning, baselining, automation)

Configuration management (CM) is a set of activities and processes performed by the change management program of an organization to ensure system configurations are consistent and secure. In this objective we will discuss topics directly related to CM, including provisioning initial configuration, maintaining baseline configurations, and automating the CM process.

Configuration Management Activities

Configuration management is part of the larger change management process. Change management is concerned with overarching strategic changes to the planning and design of the infrastructure, as well as the operational level of managing the infrastructure, while configuration management is more focused at the system level and is usually part of the tactical or day-to-day activities. Configuration management covers initial provisioning and ongoing baseline configurations of systems, applications, and other components, especially through automated processes whenever possible.

Cross-Reference

Configuration management is very closely related to patch and vulnerability management, covered in Objective 7.8, and change management, discussed in detail in Objective 7.9.

Provisioning

Just as user accounts are provisioned, as discussed in Objective 5.5, systems are also provisioned. In this context, however, provisioning is the initial installation and configuration of a system. Provisioning may require manual installation of operating systems and applications, as well as changing configuration settings to make sure that the system is both functional and secure. However, as discussed a bit later in this objective, automation can make provisioning a system far more efficient and ensure that the configuration meets its initial required baseline (discussed next).

Provisioning often uses baseline *images*, which are preapproved configurations that meet the organization's requirements for hardware and software settings, to quickly deploy operating systems and software, cutting down the time and margin for error required to install a system.

Baselining

The default settings for most systems are unsecure and often do not meet the functional needs of the organization. Therefore, the initial default configuration and settings need to be changed to better suit the organization's functional and security requirements. *Baselining* means ensuring that the configuration of a system is set according to established organizational standards and remains that way even throughout configuration and change processes. This doesn't mean that the baseline for a system won't sometimes change; baselines often change in an organization as system functions are changed, systems are upgraded, patches are applied, and the operating environment for the organization changes. Changing baselines is part of the entire change management process and must be approached with careful planning, testing, and implementation.

An organization could have several established baselines that apply to specific hosts. For example, an organization may have a workstation baseline that applies to all end-user workstations and a separate server baseline that applies to servers. It may also have baselines for network devices, and even mobile devices. The point here is that for a given device, the organization should have a baseline design that details the versions of operating systems and applications installed on the device, as well as carefully controlled configuration settings that should be standardized across all like devices.

All baseline configurations should be documented and checked periodically. There are automated software solutions, some of which are part of an operating system, that can alert an administrator if a system deviates from the baseline. Legitimate changes to baselines could be a new piece of software or even a patch that is applied to the host; these valid changes, once tested and accepted, then become part of the updated baseline configuration. It's the nonstandard or unknown changes to the baseline that must be paid attention to, however, as these may come in the form of unauthorized changes or even malware.

Baseline configuration settings often include

- Standardized versions of operating systems and applications
- Secure configuration settings, including only allowed ports, protocols, and services
- Removal or change of default account and password settings

- Removal of unused applications and services
- Operating system and application patching



EXAM TIP You should keep in mind for the exam that baselines are critical in maintaining secure configuration of all systems in the infrastructure. Secure baselines include controlled versions of operating systems and applications, as well as their security settings. An organization may have multiple baselines, depending on the type of device in question.

Automating the Configuration Management Process

Most systems in the infrastructure are complex and have a myriad of configuration settings that must be carefully set in order to maintain their security and functionality. It would be impractical for administrators to manually set all of these configuration settings and expect to have time for any other of their daily tasks. Additionally, misconfigurations occur due to human error, and many of these configuration settings may render a system nonfunctional or not secure if not set correctly. This is where automation is fundamental in maintaining configuration baselines.

Hundreds of automated tools are available to assist with configuration management. Some of them are built into the operating system itself, and many others are free or third-party utilities that come with management software. For example, Windows Server includes Active Directory (AD), which has group policy settings that enable change management administrators to manage configuration baselines in an AD domain. Linux has its own built-in configuration management utilities as well. Additionally, organizations can use powerful customized scripts, such as those written in PowerShell or Python, as well as enterprise-level management systems. We will discuss many of the tools used to configure and maintain security settings in Objective 8.2. Using automated tools to perform configuration management can help reduce issues caused by human error, ensure standardization of configuration settings across the enterprise, and make configuration changes much more efficient.

Cross-Reference

Tool sets, software configuration management, and security orchestration, automation, and response (SOAR) are related to automating the configuration management process and are discussed in depth in Objective 8.2.

REVIEW

Objective 7.3: Perform Configuration Management (CM) (e.g., provisioning, baselining, automation) This objective reviewed configuration management processes. Configuration management is a subset of change management and is closely related to both vulnerability management and patch management. The provisioning process is where the initial

installation and configuration of systems and applications occur. It's important to establish a standardized baseline to use for devices across the organization, and there may be multiple baselines to address different types of devices. Baselines also change occasionally, as the environment changes or systems and applications change. Configuration management is made much more efficient and easier by using automated tools that can help reduce human error and ensure configuration baselines are maintained.

7.3 QUESTIONS

1. Your company is creating a secure baseline for its end-user workstations. The workstations should only be able to communicate with specific applications and hosts on the network. Which of the following should be included in the secure baseline for the workstations to ensure enforcement of these restrictive communications requirements?
   - **A.** Operating system version
   - **B.** Application version
   - **C.** Limited open ports, protocols, and services
   - **D.** Default passwords
2. Riley has been manually provisioning several hosts for a secure subnet that will process sensitive data in the company. These systems are scanned before being taken out of the test environment and connected to the production network. The scans indicate a wide variety of differences in configuration settings for the hosts that have been manually provisioned. Which of the following should Riley do so that the configuration settings will be consistent and follow the secure baseline?
   - **A.** Provision the systems using automated means, such as baseline images
   - **B.** Manually configure the systems using vendor-supplied recommendations
   - **C.** Back up a generic system on a network and restore the backup to the new systems so they will be configured identically
   - **D.** Manually configure the systems using a secure baseline checklist

7.3 ANSWERS

1. **C** Any open ports, protocols, and services affect how the workstation communicates with other applications on the network or other hosts. These should be carefully considered and controlled for the secure baseline. The other choices are also considerations for the secure baseline, but do not necessarily affect communicating with only specific applications or hosts on the network.
2. **A** Riley should use an automated means to provision the secure hosts; an OS image with a secure baseline could be deployed to make the job much easier and more efficient and ensure that the configuration settings are standardized.



Objective 7.4 Apply foundational security operations concepts

In this objective we reexamine some foundational security concepts that we covered in previous domains, albeit in this objective from an operations context. These concepts include need-to-know, least privilege, separation of duties, privileged account management, job rotation, and service level agreements.

Security Operations

Security operations describes the day-to-day running of the security functions and programs. When you first learned about security theories, models, definitions, and terms, it may not have been clear as to how these things apply in the course of a security professional's normal day. Now you are going to apply the fundamental knowledge and concepts you learned earlier in the book to the operational world.

Need-to-Know/Least Privilege

Two of the important fundamental concepts introduced in Domain 1, and emphasized throughout the book, are need-to-know and the principle of least privilege. These concepts ensure that entities do not have unnecessary access to information or systems.

Need-to-Know

Recall from previous discussions that *need-to-know* means that an individual should have access only to information or systems required to perform their job functions. In other words, if their job does not require access, then they don't have the need-to-know for information, and by extension, the systems that process it. This limitation helps support and enforce the security goal of confidentiality. The need-to-know concept is applied operationally throughout security activities. Examples include restrictive permissions, rights, and privileges; the requirement for need-to-know in mandatory access control models; and the need to keep privacy information confidential.

A new employee's need-to-know should be assumed to be the minimum required to fulfill the functions of their job. As time progresses, an individual may require more access, depending on changing job requirements and the operating environment. Only then should additional access be granted. Need-to-know should be carefully considered and approved by someone with the authority to do so; normally that might mean the individual's supervisor, a data or system owner, or a senior manager. Need-to-know should also be periodically reviewed to see if the individual still has validated requirements to access systems and information. If the job requirements change or the operating environment no longer requires the individual to have the need-to-know, then access should be revoked or reduced.

Principle of Least Privilege

The principle of least privilege, as we have discussed in other objectives, essentially means that an individual should only have the rights, permissions, privileges, and access to systems and information that they need to perform their job. This may sound similar to the concept of need-to-know, but there is a subtle difference that you must be aware of for the exam. With need-to-know, an individual may or may not have access at all to a system or information. The principle of least privilege states that if an individual does have access to system or information, they can only perform certain actions. So, it becomes a matter of no access at all (need-to-know) or minimal access necessary (least privilege).



EXAM TIP Need-to-know determines *what* you can access. Least privilege regulates *what you can do* when you have access.

The principle of least privilege is applied at the operational level by only allowing individuals, ranging from normal users to administrators and executives, to perform tasks at the minimal level of permissions necessary. For example, an ordinary user should not be able to perform privileged administrative tasks on a workstation. Even a senior executive should not be able to perform those tasks since they do not relate to their duties.

Separation of Duties and Responsibilities

The concept of separation of duties (SoD) prevents a single individual from performing a critical function that may cause damage to the organization. In practice, this means that an individual should perform only certain activities, but not others that may involve a conflict of interest or allow one person to have too much power. For example, an administrator should not be able to audit their own activities, since they could conceivably delete any audit trails that record any evidence of wrongdoing. A separate auditor should be checking the activities of administrators.

Related to the concept of separation of duties are the concepts of *multiperson control* and *m-of-n control*, which require more than one person to work in concert to perform a critical task.

Multiperson Control

Multiperson control means that performing an action or task requires more than one person acting jointly. It doesn't necessarily imply that the individuals have the same or different privileges, just that the action or task requires multiple people to perform it, for the sake of checks and balances.

A classic example of multiperson control is when an individual bank teller signs a check for over a certain amount of money, and then a manager or supervisor must countersign the check authorizing the transaction. In this manner, no single individual can use this method to steal

a large amount of funds. A bank teller and bank manager could secretly agree to commit the crime, known as *collusion*, but it may be less likely because the odds of getting caught increase. Another example would be a situation that requires three people to witness and sign off on the destruction of sensitive media. One person alone can't be empowered to do this, since assigning only one person to be responsible for destroying the media could allow that person to steal the media and claim that they destroyed it. But assigning three people to witness the destruction of sensitive media would reduce the possibility of collusion and reduce the risk that the media was destroyed improperly or accessed by unauthorized individuals.

M-of-N Control

M-of-n control is the same as multiperson control, except it doesn't require all designated individuals to be present to perform a task. There may be a given number of people, "n," that have the ability to perform a task, but only so many of them (the "m") are required out of that number. For example, a secure piece of software may designate that five people are allowed to override a critical financial transaction, but only three of the five are necessary for the override to take place. This means that any three of the five people could input their credentials signifying that they agree to an override for it to take place. This doesn't necessarily imply that they all have different rights, privileges, or permissions (although in the practical world, that is often the case); it could simply mean that a single person alone can't make that decision.

Note that separation of duties does not require multiperson control or m-of-n control. An individual can have separate duties and perform those tasks daily without having to work with anyone. Multiperson or m-of-n control only comes into play when a task must be completed by multiple people working together at that moment or in a defined sequence to make a decision or complete a sensitive or critical task.

Privileged Account Management

Privileged accounts require special care and attention. In addition to carefully vetting individuals before assigning them higher-level accounts with special privileges, those accounts should be approved by the management chain. The individual should have the proper need-to-know and security clearance for a privileged account, as well as additional training that emphasizes the special procedures for safeguarding the account and the potential dire security consequences of failing to do so. Note that it's not only administrators who receive accounts with higher-level privileges; sometimes users receive accounts with additional privileges for legitimate business reasons.

Once a privileged account has been granted to an individual, they should also be carefully scrutinized for the correct authorizations. Even with a privileged account, the principles of separation of duties and least privilege still apply. Having a privileged account is not an all-or-nothing prospect; the account should still have only the privileges necessary to perform the functions related to the individual's job description, and no more. Having a privileged account

also does not mean that the individual has access to all resources and objects. There still may be sensitive data that they do not need to access, which falls under the principle of need-to-know, as mentioned earlier.



EXAM TIP Even privileged accounts are still subject to the principle of least privilege; not every privileged account requires full administrator privileges over the system or application. Privileged accounts can still be assigned only the limited rights, privileges, and permissions required to perform specific functions.

Individuals with privileged accounts should only use those accounts for specific privileged functions, and for only a limited amount of time. They should not be constantly logged into the privileged account, since that increases the attack surface for the account and the resources they are accessing. Privileged account holders should also maintain a routine user account and use it for the majority of their duties, especially for mundane tasks such as e-mail, Internet access, and so on. Using the methods described in Objective 5.2, the organization should employ just-in-time authorization; that is to say, the privileged account should only be used when and if necessary, and then the individual should revert to their basic account.

Cross-Reference

Just-in-time identification and authorization was discussed in Objective 5.2 and described the use of utilities such as `sudo` and `runas` to affect temporary privileged account access.

Privileged account management also lends itself to role-based authorization. Rather than granting additional privileges to a user account, security administrators can place the user in a role that allows additional privileges. Their membership in that role group should require that their account be audited more frequently and to a greater level of detail. This approach might be a better alternative than granting a user a separate privileged account if the majority of that user's daily job requirements necessitate use of the additional privileges. Again, the key here is frequent reviews and management approval for any additional privileges and system or information access.

Job Rotation

An organization with a job rotation policy rotates employees periodically through various positions so that a single individual is not in a position sufficiently long to conduct fraud or other malicious acts to a degree that could substantially impair the organization's ability to continue operations. Job rotation serves not only as a detective control but also as a deterrent control, because employees know that someone else will be filling their job role after a certain period of time and will be able to discover any wrongdoing. Implementing a job rotation policy in larger organizations usually is easier than in smaller organizations, which may lack multiple

people with the necessary qualifications to perform the job. Even when there is no suspicion of malicious acts or policy violations, it can be difficult to rotate someone out of a job position for normal professional growth and development since they may be so ingrained in that role that no one else can do their job. This is why planned, periodic cross-training and leveraging multiple people to understand exactly what is involved with a particular job requirement is necessary. The organization should never depend on one person only to perform a job function; this would make it very difficult to rotate a person out of that position in the event they were suspected of fraud, theft, complacency, incompetency, or other negative behaviors, let alone the critical need for having someone trained in a position for business continuity in the event an individual came to harm or departed the organization for some reason.

Mandatory Vacations

Somewhat related to job rotation is the principle of mandatory vacations—forcing an individual to take leave from a job position or even the organization for a short period of time. Frequently, if an individual simply has been performing the job function for a long period of time without a break, company policy may require that they take vacation time for rest and rehabilitation. Usually, this part of the policy allows an individual to be away from the organization for an allotted number of vacation days annually, whenever they choose. This is likely one of the more positive aspects of a mandatory vacation policy.

However, a mandatory vacation policy can also be used to force someone who is suspected of malicious acts to step away from the job position temporarily so an investigation can occur. You will often see this type of action in the news if someone in a position of public trust, for example, is suspected of wrongdoing. People are often placed on "administrative leave," with or without pay, pending an investigation. This is the same thing as a mandatory vacation. The individual may be allowed to return to their duties after the investigation completes, or they may be reassigned or even terminated from the organization.



CAUTION Organizations that actively use a mandatory vacation policy should also have, by necessity, some level of cross-training or job rotation as part of their policy and procedures.

Service Level Agreements

A service level agreement (SLA) exists between a third-party service provider and the organization. Third parties offer services, such as infrastructure, data management, maintenance, and a variety of other services, that can be provided to the organization under a contract. Note that third-party services can also include those offered by cloud providers. Service level agreements impact the security posture of an organization by affecting the security goal of availability, more often than not, but can also affect data confidentiality if a third-party has access

to sensitive information. Poor performance of services can impact the organization's efficient operations, performance, and security, so it's important to have agreements in place that ensure consistent levels of function and performance for those third-party provided services.

In addition to the legal contract documentation that will likely be included in a contract with a third-party service provider, the SLA is critical in specifying the responsibilities of both parties. This document is used to protect both the organization and the third-party service provider. The SLA can be used to guarantee specific levels of performance and function on the part of the third-party service provider, as well as delineate the security responsibilities between the customer and the provider with regard to protecting systems and data. Failing to meet SLA requirements often incurs a financial penalty.

REVIEW

Objective 7.4: Apply foundational security operations concepts In this objective we reviewed several foundational security operations concepts, including need-to-know, least privilege, separation of duties, privileged account management, job rotation, and service level agreements. Each of these concepts has been discussed in at least one previous objective, but here we framed them in the context of security operations.

Need-to-know means that an individual does not have any access to systems or information unless their job requires that access. Contrast this to the principle of least privilege, which means that once granted access to a system or information, an individual should only be allowed to perform the minimal tasks necessary to fulfill their job responsibilities. Separation of duties means that one individual should not be able to perform all the duties required to complete a critical task, thereby preventing fraudulent or malicious activity absent the collusion of two more people. This is also further demonstrated by the concepts of multiperson control and m-of-n control, which require at least a minimum number of designated, authorized individuals present to approve or perform a critical task.

Privileged account management requires that any individual having privileges above a normal user level should be vetted and approved by management for those privileges. Privileged accounts granted to these individuals should not be used for routine user functions, but only for the privileged functions they were created to perform. Privileged accounts should also be reviewed periodically to ensure they are still valid.

Job rotation is used to replace an individual in a job function periodically so that the person's activities can be audited for any malicious or wrongful acts. This is similar to mandatory vacations, which is only temporary and usually implemented while an individual is under investigation.

Service level agreements are used to protect both a third-party service provider and the organization by specifying the required performance and function levels in the contract, including security, for each party.

7.4 QUESTIONS

1. Which of the following is the best example of implementing need-to-know in an organization?
   - **A.** Denying an individual access to a shared folder of sensitive information because the individual does not have job duties that require the access
   - **B.** Allowing an individual to have read permissions, but not write permissions, to a shared folder containing sensitive information
   - **C.** Requiring the concurrence of three people out of four who are authorized to approve a deletion of audit logs
   - **D.** Routinely reassigning personnel to different security positions that each require access to different sensitive information
2. Audit trails for a sensitive system have been deleted. Only a few people in the company have the level of training and privileged access required to perform that action. Although a particular person is suspected of performing the malicious act, all people who have access must be removed from their position, at least on a temporary basis, during the investigation. Which of the following does this action describe?
   - **A.** Separation of duties
   - **B.** Job rotation
   - **C.** Mandatory vacation
   - **D.** M-of-n control

7.4 ANSWERS

1. **A** Need-to-know is typically a deny or allow situation; denying access to a shared folder containing sensitive information that the user does not require for their job duties is based on need-to-know.
2. **C** Since only a few people have that level of access, they must all be temporarily removed from their positions during the investigation and placed on administrative leave, a form of mandatory vacation. Job rotation is not an option if there are only a few people who can perform the job function and they are all under investigation.



Objective 7.5 Apply resource protection

We have discussed protecting resources throughout this entire book, but in this objective we're going to focus specifically on one area we have not previously addressed—media management and protection. Media is often associated with backup tapes, but it also includes

hard drive arrays, CD-ROM and Blu-ray discs, storage area networks (SANs), network-attached storage (NAS), and portable media such as USB thumb drives, regardless of whether they are local or remote storage. This objective focuses on managing the wide variety of media and the specific security measures used to protect it.

Media Management and Protection

Regardless of the type of media your organization is using, you should carefully consider several key protection activities related to media management. These include administrative, technical, and physical security controls. Each type of control is applied to protect media from unauthorized access, theft, damage, and destruction. We'll discuss some of these controls in the upcoming sections.

Media Management

Media management primarily uses administrative controls, such as policies and procedures, associated with dictating how media will be used in the organization. Management should create a media protection and use policy that outlines the requirements for proper care and use of storage media in the organization. This policy could also be closely tied to the organization's data sensitivity policy, in that the data residing on media should be protected at the highest level of sensitivity dictated by the policy.

Media management requirements detailed in the policy should include

- All media must be maintained under inventory control procedures and secured during storage, transportation, and use.
- Proper access controls, such as object permissions, must be assigned to media.
- Only authorized portable media should be used in organizational systems, and portable media must be encrypted.
- Media should only be reused if sensitive data can be adequately wiped from it.
- Media should be considered for destruction if it cannot be reused due to the sensitivity of data stored on it.

Media Protection Techniques

Media management sets administrative policy controls for the use, storage, transport, and disposal of various types of media in the organization. It also dictates the practical controls expected for those activities. Protection techniques for media include technical and physical controls implemented during access, transport, storage, and disposal.

Media Access Controls

Media should be treated with care and handling commensurate with the level of sensitivity of data stored on the media. This again should be in accordance with the data sensitivity policies determined by management. These controls include

- Sensitive data stored on media should be encrypted.
- Access control permissions granted to authorized users should be based upon their job duties and need-to-know; the principles of least privilege and separation of duties should also be included in these access controls.
- Strong authentication mechanisms required to access media must be used even for authorized users.

Media Storage and Transportation

Physical controls are the primary type of control used to secure media while it is stored and during transport to prevent access by unauthorized personnel. Physical controls for media storage and transport must include

- Secure media storage areas (e.g., locked closets and rooms)
- Physical access control lists of personnel authorized to enter media storage areas
- Proper temperature and humidity controls for media storage locations
- Media inventory and accountability systems
- Proper labeling of all media, including point-of-contact information (i.e., data or system owner), sensitivity level, archival or backup date, and any special handling instructions
- Two-person integrity for transporting media containing highly sensitive information (requiring two people to witness/control the transportation of highly sensitive media for security)



EXAM TIP Key media protection controls include media usage policies, data encryption, strong authentication methods, and physical protection.

Media Sanitization and Destruction

Media should be kept only as long as it is needed. Recall that Objective 2.4 explained why data should not be retained longer than the organization requires it for legitimate business reasons or due to regulatory requirements. Once information is not needed for either reason, the organization should dispose of it in accordance with policy. This includes any media that

contains the data. Often media can be sanitized or cleared for reuse within the organization, but in certain circumstances it must be destroyed to prevent any chance that sensitive information may inadvertently fall into the wrong hands.

Sanitization methods should be more thorough than simply formatting or repartitioning media; data can easily be recovered through forensic processes or by using common file recovery tools. Wiping is a much better way to clear media that is intended for reuse. *Wiping* involves writing set patterns of ones and zeros to the media to overwrite any remnant data that may still exist on the media even after file deletion or media formatting.

Media destruction should be used when media is worn out, obsolete, or otherwise not intended to be reused. In cases where highly sensitive data resides on the media, wiping techniques may not be enough to give management the confidence that data may not be recovered by someone using advanced forensic tools and techniques. In these cases, media should simply be destroyed. Media destruction methods include degaussing, burning, pulverizing, and even physical destruction using hammers or other methods to break the media, rendering it completely unusable. For highly sensitive media, the organization should consider implementing two-person integrity; that is to say that it requires two people to participate in and witness the destruction of sensitive media so that management can be assured it will not fall into the wrong hands.

Cross-Reference

Data retention, remanence, and destruction were also discussed in Objective 2.4.

REVIEW

Objective 7.5: Apply resource protection This objective examined resource protection, specifically focusing on media management and the controls used to protect the variety of media types. Media management begins with policies, which should include inventory control, access control, and physical protections. Media protection controls include those implemented to protect media during use, storage, transportation, and disposal. Specific controls include the need for encryption, strong authentication, and object access. Media must be sanitized to erase any sensitive data remnants if it will be reused. If reusing media is not practical, it must be destroyed.

7.5 QUESTIONS

1. Which of the following must media management and protection begin with?
   - **A.** Media policies
   - **B.** Strong encryption
   - **C.** Strong authentication
   - **D.** Physical protections

2. Management has made the decision to destroy media that contains sensitive data, rather than reuse it. Because this media might fetch a good price from the organization's competitors, management wants to put in place additional controls to make sure that the media is destroyed properly. Which one of the following would be an effective control during media destruction?
   - **A.** Destruction documentation
   - **B.** Burning or degaussing media
   - **C.** Two-person integrity
   - **D.** Strong encryption mechanisms

7.5 ANSWERS

1. **A** Media protection begins with a comprehensive media use policy, established by organizational management, which dictates the requirements for media use, storage, transportation, and disposal.
2. **C** In this scenario, one of the most effective security controls to ensure that media has been destroyed properly is the use of a two-person integrity system, which requires two people to participate in and witness the destruction of sensitive media so that management can be assured it will not fall into the wrong hands.



Objective 7.6 Conduct incident management

In this objective we will cover the phases of the incident management process that you need to know for the CISSP exam, which include detection, response, mitigation, reporting, recovery, remediation, and lessons learned. We'll also look at another phase, preparation, that is commonly identified as the first phase in many other incident management life cycles. Keep in mind as you read this objective that incident management differs from disaster recovery planning and processes (covered in Objectives 7.11 and 7.12) and business continuity planning (covered in Objective 7.13), though they sometimes overlap depending on the nature of the incident.

Security Incident Management

An *incident* is any type of event with negative consequences for the organization. As cybersecurity professionals, we often categorize incidents as some form of infrastructure attack, but even a temporary power outage, server failure, or human error technically falls into the category of incidents if it negatively impacts the organization. Any incident which affects the security of the organization or its assets, whether it stems from a malicious hacker, a natural disaster, or

simply the action of a complacent employee, is of concern to cybersecurity professionals. Any event that affects the three security goals of confidentiality, integrity, and availability could be considered a security incident. Incident management is much more than simply responding to an incident. The entire incident management process includes a management program and process used by the organization to plan and execute preparations for an incident, respond during the incident, and conduct post-incident activities.



EXAM TIP A security incident doesn't necessarily involve a malicious act; it can also be the result of a natural disaster such as a flood or tornado, or from an accident such as a fire. It can also be the result of a negligent employee.

Incident Management Life Cycle

As a formalized process, incident management has a defined life cycle. A variety of incident management life cycles are promulgated by different books, standards, and professional organizations. Although their titles and specific phases usually differ, they all promote the phases of incident detection, response, and post-incident activities. As one example, the National Institute of Standards and Technology (NIST) offers an incident response life cycle model in its Special Publication (SP) 800-61, Revision 2, *Computer Security Incident Handling Guide*, which discusses the life cycle phases of preparation; detection and analysis; containment, eradication, and recovery; and post-incident activity. Figure 7.6-1 illustrates the NIST life cycle.

CISSP exam objective 7.6 outlines similar phases that could easily be mapped to the NIST life cycle model or any one of several other incident management life cycle models. The point here is that incident management is a formalized process that should not be left to chance.

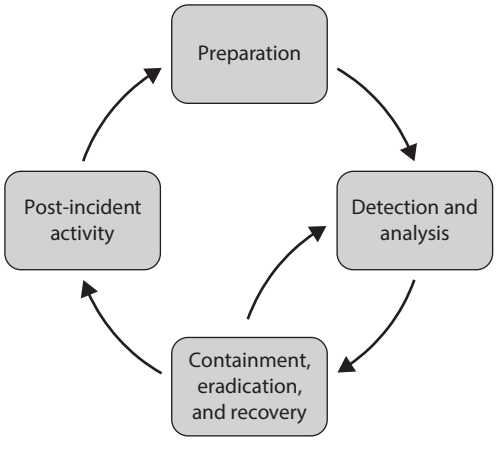


FIGURE 7.6-1 The NIST incident response life cycle (adapted from NIST SP 800-61 Rev. 2, Figure 3-1)

Every organization should have formal incident management policies and procedures, an adopted standard for incident response, and a formal incident management life cycle that the organization adheres to.

Preparation

Oddly enough, the preparation phase is not part of the formal CISSP exam objectives for incident management; however, it is still an important concept you should be familiar with since the other steps of the incident management process rely so much on adequate preparation. The NIST life cycle model discusses this phase as being critical in overall incident management and describes *preparation* as having all the correct processes in place, as well as the supporting procedures, equipment, personnel, information, and other needed resources.

The preparation phase of incident management includes

- Development of the incident management strategy, policy, plan, and procedures
- Staffing and training a qualified incident response team
- Providing facilities, equipment, and supplies for the incident response capability

The procedures that an organization must develop for incident management come from incident response policy requirements and must take into account the potential need for different processes during an incident than the organization normally follows day to day. These processes should be tailored around incident management and include

- Incident definition and classification
- Incident communications procedures, including notification and alerting, escalation, reporting, and communications with both internal stakeholders and external agencies
- Incident triage, prioritization, and escalation
- Preservation of evidence and other forensic procedures
- Incident analysis
- Attack containment
- Recovery procedures
- Damage mitigation and remediation

Detection

Most detection capabilities are not focused on incident response, but rather on incident prevention. Detection capabilities should be included as a normal part of the infrastructure architecture and design. These capabilities include intrusion detection and prevention systems (IDS/IPSs), alarms, auditing mechanisms, and so on.

Early detection is one of the most important factors in responding to an incident. Detection mechanisms must be tuned appropriately so that they catch seemingly unimportant singular events that may indicate an attack (called *indicators of compromise*, or *IoCs*) but are not prone

to reporting false positives. This is a very delicate balance, and one that will never be completely perfect. As the organization matures its incident management capability, the number of false positives will decrease, allowing the organization to identify patterns that indicate an actual incident.

Detection is based on data that comes from a variety of sources, including anti-malware applications, device and application logs, intrusion detection alerts, and even situational awareness of end users who may report anomalies in using the network.

Response

Once an incident is detected, there are several things that must occur quickly. First, the incident must be triaged to determine if it is a false positive and, if not, determine its scope and potential critical impact. Organizations often develop checklists that IT, security, and even end-user personnel can use to determine if an incident exists and, if so, how serious it is and what to do next. For end users, this list is usually very basic and ends up with the correct action being to report the incident to security personnel. For IT and cybersecurity employees, this checklist will be much more involved, with multiple possible decision points.

When the incident is appropriately triaged, the incident response team is notified, and, if necessary, the incident is escalated to upper management or outside agencies. If the incident is considered any type of disaster, particularly one which could threaten human safety or cause serious damage to facilities or equipment, the disaster response team is also notified. Usually the decision to notify outside agencies must come from a senior manager with authority to make that decision. The response phase is also when the incident response team is activated. Very often this notification comes from a 24/7 security operations center (SOC), on-call person, or an incident response team member. A call tree is often activated to ensure that team members get notified quickly and effectively. In some cases, the incident response command center, if it is not already part of the SOC, is activated. Each team member has a job to do and transitions from their normal day-to-day position to their incident-handling jobs.

The incident response (IR) team has several key tasks it must perform quickly. Almost simultaneously, the IR team must gather data and analyze the cause of the incident, the scope, what parts of the infrastructure the incident is affecting, and which systems and data are affected. The IR team must also work quickly to contain the incident as soon as possible, to prevent its spread, and find the source of the incident and stop or eradicate it. All these simultaneous actions make for a very complex response, especially in large environments.

In addition to analysis, containment, and eradication, the IR team must also make every attempt to gather and preserve forensic evidence necessary to determine what happened and trace the incident back to its root cause. Evidence is also necessary to ensure that the responsible parties are discovered and held accountable.

Cross-Reference

Investigations and forensic procedures were discussed in Objective 7.1.

The initial response is not considered complete until the incident is contained and halted. For instance, a malware spread must be stopped from further damaging systems and data, a hacking attack must be blocked, and even a non-malicious incident, such as a server room flood, must be stopped. Once the incident has been contained and the source prevented from doing any further harm, the organization must now turn its attention to restoring system function and data so the business processes can resume.

Mitigation

Mitigating damage during the response has many facets. First, the incident must be contained and the spread of any damage must be limited as much as possible. Sometimes this requires implementing temporary measures. These can range from temporarily shutting down systems, rerouting networks, and halting processing to more drastic steps. But these are only temporary mitigations necessary to contain the incident; permanent mitigations may also have to be considered, sometimes even while the incident is still occurring. Temporary corrective controls like emergency patches, configuration changes, or restoring data from backups will often be put in place until more permanent solutions can be implemented. Permanent or long-term mitigations are covered later during the discussion on remediation.



EXAM TIP Remember that corrective controls are temporary in nature and are put in place to immediately mitigate a serious security issue, such as those that occur during an incident. Compensating controls are longer-term in nature and are put in place when a preferred control cannot be implemented for some reason. The difference between corrective and compensating controls was also discussed in Objective 1.10.

Reporting

There are many aspects to reporting, both during and after an incident. Effective reporting is highly dependent on the communications procedures established in the incident response plan. During the incident, reports of the status of the response, especially efforts to contain and eradicate the incident, are communicated up and down the chain of command, as well as laterally within the organization to other departments affected by the incident. During and after more serious incidents, reporting to external third parties may occur, such as law enforcement, customers, business partners, and other stakeholders. Reporting during the incident may occur several times a day and may be informal or formal communications such as status e-mails, phone calls, press conferences, or even summary reports at the end of the response day.

The other facet of reporting is post-incident reporting, which requires more formal and comprehensive reports. Note that post-incident reporting normally takes place after the remediation step is completed, as discussed later on in this objective. Reports must be delivered to key stakeholders both within the organization and outside it. Senior management must decide which

sensitive information should be reported to various stakeholders, since some of the information may be proprietary or confidential. In any event, the incident response team develops a report that summarizes the incident for nontechnical people, but it may have technical appendices. The report includes the root cause analysis of the incident, the responses actions, the timeframe of the incident, and what mitigations were put in place to contain and eradicate the cause. The report also usually includes recommendations to prevent further incidents.

Recovery

Recovery efforts take place after an incident has been contained and the cause mitigated. During this phase of the incident management process, systems and data are restored and the business operations are brought back online. The goal is to bring the business back to a fully operational state as soon as possible, but that does not always happen if the damage is too extensive. If systems have been damaged or data is lost, the organization may operate in a degraded state for some time.

This phase of incident management tests the effectiveness of the organization's business continuity planning, if the incident is serious enough to disrupt business operations. This is one point where incident response is directly related to business continuity. During the business continuity planning process, the business impact analysis defines the critical business processes and the systems and information that support them, so they can be prioritized for restoration after an incident.

Cross-Reference

Business impact analysis and business continuity were both discussed in Objective 1.8, and business continuity will be discussed in depth in Objective 7.13.

Remediation

Remediation addresses the long-term mitigations that repair the damage to the infrastructure, including replacement of lost systems and recovery of data, as well as implementation of solutions to prevent future incidents of the same type. The organization must develop a plan to remediate issues that caused the incident, including any vulnerabilities, lack of resources, deficiencies in the security program, management issues, and so on. At this point, the organization should perform an updated risk assessment and analysis. This allows the organization to reassess its risk and see if it failed to implement sufficient risk reduction measures, as well as identify new risks or update the severity rating of previously known ones.

This phase of the incident management life cycle is just as much managerial as it is technical. Vulnerabilities can be patched and systems can be rebuilt, but management failures are often found to be the root causes of incidents. Management must recommit to providing needed resources, such as money, people, equipment, facilities, and so on. This is all part of the remediation process.

Lessons Learned

The final piece of incident management is understanding and implementing lessons learned from the incident response. The organization must perform in-depth analysis to determine why the incident occurred, what could have prevented it, and what must be done in the future to prevent a similar incident from occurring again. Lessons learned should be included in the final report, but they must also be ingrained in the organization's culture so that these lessons can be used to protect organizational assets from further incident.

Lessons learned don't have to be limited to looking at the organization's failures that may have led up to the incident; they can also look at how the organization planned, implemented, and executed its incident response. Some of these lessons learned may include ways to improve the following:

- Response time, including incident detection, notification, escalation, and response
- Deployment of resources during an incident, including people, equipment, and time
- Personnel staffing and training
- The incident response policy, plan, and procedures

In any event, examining the entire incident management life cycle for the organization after a response will glean many lessons that the organization can use in the future, provided it is willing to do so.

REVIEW

Objective 7.6: Conduct incident management In this objective we examined the incident management program within an organization. We reviewed the need to adopt an incident management life cycle, of which there are many, and briefly examined one in particular promulgated by NIST. We then discussed the various phases of the incident management process that you will need to understand for the CISSP exam.

- Preparation is the most important phase of the incident management process, since the remainder of the response depends on how well the organization has prepared itself for incidents.
- Early detection of an incident is extremely critical so that the organization can execute its response rapidly and efficiently.
- The response itself has many pieces to it, including incident containment, analysis, and eradication of the cause of the incident.
- The mitigation phase consists of implementing temporary measures, in the form of corrective controls that can preserve systems, data, and equipment and keep the business functioning at some level; but corrective controls need to be replaced with more permanent and carefully considered mitigations during the remediation phase.

- Reporting includes all the communications that are necessary both during and after the incident. Reporting can include communications up and down the chain of command, as well as laterally across the organization. An effective communications process should be included in the incident response plan. It may also require reporting to parties outside the organization, such as law enforcement, regulatory agencies, or partners and customers. A formal report should be generated after the incident that includes a comprehensive analysis of the root cause and recommendations for preventing further incidents.
- The incident recovery phase involves bringing the business back to a fully operational state after an incident, which may take time and happen in phases depending upon how serious the impact of the incident has been. Recovery operations include the prioritized restoration of systems and data based on a thorough business impact analysis, which is performed during the business continuity planning process.
- Remediation after an incident consists of the more permanent controls that must be implemented to repair damage to systems and prevent the incident from recurring.
- Understanding lessons learned requires examining the entire incident management process to determine deficiencies in the organization's security posture, as well as its incident response processes. These lessons must be understood and used to protect the infrastructure from further incidents.

7.6 QUESTIONS

1. During which phase of the incident management life cycle is the incident response plan developed and the incident response team staffed and trained?
   - **A.** Preparation
   - **B.** Response
   - **C.** Lessons learned
   - **D.** Recovery
2. Your organization is in the early stages of responding to an incident in which malware has infiltrated the infrastructure and is rapidly spreading across the network, systematically rendering systems unusable and deleting data. Which of the following actions is one of the most critical in stopping the spread of the malware to prevent further damage?
   - **A.** Analysis
   - **B.** Triage
   - **C.** Escalation
   - **D.** Containment

7.6 ANSWERS

1. **A** All planning for incident response, including developing the actual incident response plan and fielding the response team, is conducted during the preparation phase of the incident management life cycle. Performing these activities during any of the other phases of the incident management life cycle would be too late and largely ineffective.
2. **D** Containment is likely the most critical activity an incident response team should engage in since this prevents further damage to systems and data. The other answers are also important but may not directly contribute to stopping the spread of the malware.



Objective 7.7 Operate and maintain detective and preventative measures

In this objective we focus on technical controls that are considered preventive and detective in nature. Prevention is preferred so that negative activity can be stopped before it even begins; however, absent prevention, rapid detection is critical to quickly stopping an incident to contain and minimize its damage to the infrastructure. We will discuss firewalls and intrusion detection/prevention systems and how they work. We will also briefly explore third-party services and their role in security. In addition, we will examine various other preventive and detective controls used, such as sandboxing, honeypots and honeynets, and the all-important and ubiquitous anti-malware controls. Finally, we will discuss the roles that machine learning and artificial intelligence play in cybersecurity.

Cross-Reference

Control types and functions were discussed in Objective 1.10.

Detective and Preventive Controls

Of all the control functions we have discussed throughout the book, prevention is arguably the most important. Preventing an incident from occurring is very desirable in that it saves time, money, and other resources, as well as prevents loss of or damage to information assets. We have available a multitude of different administrative, technical, and physical controls that are focused on preventing illegal acts, violations of policy, emergency or disaster situations, data loss, and so on. However, preventive controls are not always enough. Despite having a well-designed and architected security infrastructure that uses defense in depth as a secure design principle, incidents still happen. Even before incidents occur, detective controls must be in place since early detection of an incident can help to reduce the level of damage done to systems, data, facilities, equipment, the organization, and most importantly, people.

Allow-Listing and Deny-Listing

Allow-listing and deny-listing (formally known as *whitelisting* and *blacklisting*, respectively) are techniques that can allow or deny (block) items in a list. These listings are essentially rule sets. A *rule set* is a collection of rules that stipulate allow and deny actions based on specific content items, such as types of network traffic, file types, and access control list entries. These rule sets are used to control what is processed, transmitted or received, or accessed in an infrastructure.

Most rule sets depend on context, since they can be used in different ways in security. For example, a rule set that lists allowed applications or denied applications can be used, respectively, to allow the corresponding executables to run on a system or deny them from running on a system. Another use is allowing or denying certain types of network traffic based on characteristics of that traffic, such as port, protocol, service, source or destination host IP address, domains, and so on. Still another implementation of rule sets might allow or deny access to a resource, such as a shared folder, by individuals or groups of users.



EXAM TIP As with everything in technology, concepts and terms change from time to time, based on newer technologies, the environment we live and work in, and even social change. And so it goes for the terms *whitelist* and *blacklist*, which have been deprecated and are decreasing in use within our professional security community. In fact, (ISC)² indicates in their own blog post (https://blog.isc2.org/isc2_blog/2021/05/isc2-supports-nist.html) that they intend to follow NIST's lead to discontinue the terms "blacklisting" and "whitelisting." In anticipation of their changes in terminology, I will use the inclusive terms *allow list* and *deny list*, respectively. However, be aware that because the CISSP exam objectives may not have caught up with this change at the time of this writing, you may still see the terms "whitelist" and "blacklist" on the exam.

Often these rule sets are implemented in *access control lists (ACLs)*, a term normally associated with network devices and traffic. While modern allow/deny lists may be combined into a single monolithic rule set that has both allow and deny entries in it, you may still see lists that exclusively allow or exclusively deny the items in the rule set. By way of explanation, here's how those exclusive lists work:

- An allow list is used to allow only the items in that rule set to be processed, transmitted, received, or accessed. Since the items in this list are the exceptions that are allowed to process, anything not on the list is, by default, denied. Although called an allow list, this is also what implements a default-deny method of controlling access, since by default everything is denied unless it is in the list.
- A deny list works the exact opposite of an allow list. All the elements of the rule set are denied. Anything not in the rule set is allowed. This is called a default-allow method of controlling access, since anything not in the list is, by default, allowed to process through the rule set.



EXAM TIP The terminology can be somewhat confusing, but an *allow* list enables a default-*deny* method of controlling access, since anything that is not in the list is not processed, and a *deny* list enables a default-*allow* method of controlling access, since anything that is not in the list is processed.

Note that, as mentioned a bit earlier, modern rule sets have entries that simply have both allow and deny rules in them, so access is carefully controlled. However, whether the organization uses as the access control method a list with both allow and deny entries in it, or a default-deny or a default-allow paradigm is often based on the organization's network resource policies regarding openness, transparency, and permissiveness. This is a good example of how an organization's appetite and tolerance for risk is connected to how it implements technical controls; an organization that has a high tolerance for risk might implement a default-allow method of access control, which is far less restrictive than a default-deny mentality.

Allow- and deny-listing is a very important fundamental concept to understand for both the real world and the CISSP exam since this technique is used throughout security. Allow and deny lists can be used separately and together on network security devices such as firewalls, intrusion detection and prevention systems, border routers, proxies, and so on. These techniques are also used to restrict software that is allowed to run on the network, as well as control which subjects can access which objects in the infrastructure.

You'll also encounter the following terms in the context of allow- and deny-listing:

- **Explicit** Refers to actual entries in an allow list or deny list. The entries in a deny list are items that are explicitly denied and the entries in an allow list are items that are explicitly allowed.
- **Implicit** Refers to anything that is not listed but, by implication, is allowed (in the case of a deny list) or denied (in the case of an allow list).

Firewalls

For better or for worse, firewalls have traditionally been considered by both security professionals and laypeople to be the ubiquitous be-all and end-all of security protection. However, firewalls do not take care of every security issue in the infrastructure. Firewalls are simply devices that are used to filter traffic from one point to another. Firewalls use rule sets as well

as other advanced methods of inspecting network traffic to make decisions about whether to allow or deny that traffic to specific parts of the infrastructure. Most firewalls are either network based or host based, but other, more recent types of firewalls also are available, including web application firewalls and cloud-based firewalls.

Network- and Host-Based Firewalls

Traditional network-based firewalls provide separation and segmentation for different parts of the network. In a traditional firewall deployment, a firewall sits on the network perimeter of an organization, separating the public Internet from the internal organizational network. Network-based firewalls may also be deployed in a *demilitarized zone (DMZ)* or *screened subnet* architecture, which makes use of other security devices, such as border routers, in combination with one or more firewalls so that traffic is not only segmented but also routed to other network segments. In a DMZ architecture, traffic enters an external firewall, is appropriately examined and filtered (allowed or blocked), and then may be redirected to another piece of the network that is not part of the internal network. Traffic inbound for the internal network may proceed through an internal firewall before it gets to its destination. This enables multiple layers of traffic filtering.

Host-based firewalls are far less complex in nature and only protect a particular host. They may be integrated with other host-based security services, such as anti-malware or intrusion detection and prevention systems (discussed in an upcoming section). Host-based firewalls are normally not dedicated security appliances; they are simply software installed as an application on the host or, in some cases, come as part of the operating system, such as Windows Defender Firewall.

Firewall Types

Although we discussed firewalls in Objective 4.2, it's helpful for CISSP exam preparation purposes to review them in the context of security operations and to introduce a few more firewall types used in security operations, such as web application and cloud-based firewalls. Network-based firewalls have more than one network interface, allowing them to span multiple physical and logical network segments, which enables them to perform traffic filtering and control functions between networks. Firewalls also use a variety of criteria to perform filtering, including traffic characteristics and patterns, such as port, protocol, service, source or destination addresses, and domain. Advanced firewalls can even filter based on the content of network traffic.

As a review of Objective 4.2, the primary types and generations of firewalls are as follows:

- *Packet-filtering* or *static firewalls* filter based on very basic traffic characteristics, such as IP address, port, or protocol. These firewalls operate primarily at the network layer of the OSI model (TCP/IP Internet layer) and are also known as *screening routers*; these are considered first generation firewalls.

- *Circuit-level firewalls* filter session layer traffic based on the end-to-end communication sessions rather than traffic content.
- *Application-layer firewalls*, also called *proxy firewalls*, filter traffic based on characteristics of applications, such as e-mail, web traffic, and so on. These firewalls are considered second-generation firewalls, which work at the application layer of the OSI model.
- *Stateful inspection firewalls*, considered third-generation firewalls, are dynamic in nature; they filter based on the connection state of the inbound and outbound network traffic. They are based on determining the state of established connections. Remember that stateful inspection firewalls work at layers 3 and 4 of the OSI model (network and transport, respectively)
- *Next-generation firewalls (NGFWs)* are typically multifunction devices that incorporate firewall, proxy, and intrusion detection/prevention services. They filter traffic based on any combination of all the techniques of other firewalls, including deep packet inspection (DPI), connection state, and basic TCP/IP characteristics. NGFWs can work at multiple layers of the OSI model, but primarily function at layer 7, the application layer.

Web Application Firewalls

A web application firewall (WAF) is a newer, special-purpose firewall type. It's used specifically to protect web application servers from web-based attacks, such as SQL and command injection, buffer overflow attacks, and cross-site scripting. WAFs can also perform a variety of other functions, including authentication and authorization services through on-premises IdM services, including Active Directory, and third-party IdM providers.

Cross-Reference

Identify management (IdM) was introduced in Objective 5.2.

Cloud-Based Firewalls

Another recent development in firewall technology involves the use of cloud-based firewalls offered by cloud service providers. As we will discuss in an upcoming section on third-party security services, many organizations do not have the qualified staff available to manage security functions within the organization, so they outsource these functions to a third-party service provider. In the case of cloud-based firewalls, a third party provides Firewall as a Service (FWaaS), which consists of managing and maintaining firewall services, normally for organizations that also use other cloud-based subscriptions, such as Platform as a Service or Infrastructure as a Service. Note that while deploying a cloud-based firewall alone can greatly simplify management of the security infrastructure for the organization, using a cloud-based firewall when a larger portion of the organization's infrastructure has migrated into the cloud makes it all the more effective.

Intrusion Detection Systems and Intrusion Prevention Systems

Historically, intrusion detection systems (IDSs) were focused on simply detecting potentially harmful events and alerting security administrators. Then, more advanced intrusion prevention systems (IPSs) were developed that could actually prevent intrusions by dynamically rerouting traffic or by making advanced filtering (allow and deny) decisions during an attack. Over the course of a few generations of technology changes, IDSs and IPSs have merged and essentially become integrated. Although an IDS/IPS could be a standalone system, typically IDS/IPS functions are part of an advanced or next-generation security system that integrates those functions, as well as firewall and proxy functions, into a single system, typically a dedicated hardware appliance or software suite.

Traditional IDS/IPSs collect and analyze traffic by forcing traffic to flow into one interface and out another, which requires the IDS/IPSs to be placed inline within the network infrastructure. The problem with this approach is that it introduces latency into the network, since the IDS/IPS's rule set must examine every packet that comes through the system. Advances in technology, however, allow an IDS/IPS to be placed at strategic points in the infrastructure, with sensors deployed across the network in a distributed environment, so that traffic is not forced to go through a single chokepoint. This reduces latency and allows the IDS/IPS to have visibility into more network segments.

IDS/IPSs are also categorized in terms of whether they are network-based or host-based:

- **Network-based IDS/IPS (NIDS/NIPS)** Looks primarily at network traffic entering into an infrastructure, exiting from the infrastructure, and traveling between internal hosts. A NIDS/NIPS does not screen traffic exiting or entering a host's network interface.
- **Host-based IDS/IPS (HIDS/HIPS)** Looks at traffic entering and exiting a specific host, and is typically implemented as software installed on the host as a separate application or as part of the operating system itself.

In addition to monitoring traffic for the host, the HIDS/HIPS may be integrated with other security software functions and may perform traffic filtering for the host, anti-malware functions, and even advanced endpoint monitoring and protection. Although not required, most modern HIDSs/HIPSs in large enterprises are agent-based, centrally managed systems. They use software endpoint agents installed on the host so security information can be reported back to a centralized collection point and analyzed individually or in aggregate by a SIEM, as discussed in Objective 7.2.

Objective 7.2 also discussed the methods by which an IDS/IPS detects anomalous network traffic and potential attacks. To recap, there are three primary methods that can be used alone or in combination to detect potential issues in the network:

- *Signature- or pattern-based (rule-based) detection* (also called *knowledge-based detection*) uses preconfigured attack signatures that may be included as part of a subscription service from the IDS/IPS vendor.
- *Anomaly- or behavior-based analysis* involves allowing the IDS/IPS to "learn" the normal network traffic patterns in the infrastructure; when the IDS/IPS detects a deviance from these normal patterns, the system alerts an administrator.
- *Heuristic analysis* takes behavior analysis one step further; in addition to catching changes in normal behavior patterns, a heuristic analysis engine looks at those abnormal behaviors to determine what types of malicious activities they could lead to on the network.



EXAM TIP You should understand the methods by which IDS/IPSs detect anomalies and potential intrusions, as well as how they are classified as either network-based or host-based systems.

Note that IDS/IPSs can look at a multitude of traffic characteristics to detect anomalies and potentially malicious activities, including port, protocol, service, source and destination addresses, domains, and so on. These characteristics could also include particular patterns like abnormally high bandwidth usage or network usage during a particular time of day or night when traffic usually is light. Advanced systems can also do in-depth content inspection of specific protocols, such as HTTP, and even intercept and break secure connections using protocols such as TLS, so that the systems can detect potentially malicious traffic that is encrypted within secure protocols.

Cross-Reference

Intrusion detection and prevention were also discussed in Objective 7.2.

Third-Party Provided Security Services

Organizations, particularly smaller ones, may not always be staffed sufficiently to take care of their own security services and infrastructure. With the mounting security challenges that organizations now face, there is an increasing trend in the use of *managed security services (MSSs)*, also known as *Security as a Service*, third-party providers to which organizations can outsource some or all of their security functions. An MSS can manage various aspects of an organization's security, such as security device configuration, maintenance, and monitoring, security operations center (SOC) services, and even user and resource management or control.

Contracting services out to a reputable third-party service provider has the following advantages, among others:

- **Cost savings** The organization does not have to hire and train its own security personnel, nor maintain a security infrastructure.
- **Risk sharing** Since the organization does not maintain its own security infrastructure, some of the risk involved with this endeavor is shared with another party.

However, there are also distinct disadvantages to contracting with a third-party service provider:

- **Less control over the infrastructure** The organization does not always have the ability to immediately control how the infrastructure is configured or react to both customer needs and events. It relies on the third party to be dependable in its responsiveness, as well as have a sense of urgency.
- **Legal liability** In the event of a breach, the organization still retains ultimate responsibility and accountability for sensitive information (although the third-party service provider may also have some degree of liability).
- **Lack of visibility into the service provider's infrastructure** The organization may not even be able to look at its own audit logs or security device performance. The organization may also not have the ability to audit the third-party security provider's processes, legal or regulatory compliance, or infrastructure.

Regardless of the security functions an organization chooses to outsource, the key to a successful relationship with a third-party security service provider is the service level agreement (SLA). An MSS usually offers a standard SLA, which the organization should carefully review and, if necessary, seek modification of before entering into the contract with the provider. The SLA should clearly define both distinct and shared responsibilities the organization and the third-party provider have in securing systems and information. The SLA should also address critical topics such as availability, resiliency, data information ownership, and legal liability in the event that an incident such as a breach occurs.

Cross-Reference

Third-party providers and some of the services they offer, as well as service level agreements, are discussed in detail throughout Objectives 1.12, 4.3, 5.3, and 8.4.

Honeypots and Honeynets

In their ongoing effort to prevent attackers from getting to sensitive hosts, administrators often deploy a honeypot on the network as part of their defense-in-depth strategy. A *honeypot* is an intentionally vulnerable host, segregated from the live network, that appears to attackers as a

prime target to exploit. The honeypot distracts the attacker from sensitive hosts, and at the same time gives the administrator an opportunity to record and review the attack methods used by the attacker.

Honeypots are often deployed as virtual machines and are segmented from sensitive hosts by both physical and virtual means. They may be on their own physical subnet off of a router, as well as use VLANs that are tightly controlled. They may have dedicated IDS/IPSs monitoring them, as well as other security devices. Administrators often have the option of dynamically changing the honeypot's configuration or disabling it altogether if needed in response to an attacker's actions.

A sophisticated attacker may recognize a lone honeypot, so a more advanced technique network defenders may deploy is a honeynet. A *honeynet* is a network of honeypots that can simulate an entire network, including infrastructure devices, servers, end-user workstations, and even security devices. The attacker may be so busy trying to navigate around and attack the honeynet that they do not have time to attack actual sensitive hosts before a security administrator detects and halts the attack.

Note that an organization should carefully consider the use of honeypots and honeynets before deciding to deploy them. If implemented improperly, a honeypot/honeynet can cause legal issues for an organization, since attackers have been known to use a honeypot to further attack a different network outside the organization's control. This could subject the organization to potential legal liability. Additionally, it can be a legal gray area if an organization tries to press charges against an attacker, as the attacker might be able to claim they were entrapped, particularly if the honeypot was set up by a law enforcement or government agency. The organization should definitely consult with its legal department before deploying honeypot technologies.

Anti-malware

Malware is a common and prevalent threat in today's world. Most organizations take malware seriously and install anti-malware products on both hosts and the enterprise infrastructure. Much of the malware that we see today is referred to as *commodity malware* (aka commercial-off-the-shelf malware). This is common malware that malicious entities obtain online (often free or cheap) to use to attack organizations. It normally targets and attacks organizations that don't do a good job of managing vulnerabilities and patches in their system, and it looks for easy targets that may not update their anti-malware software on a continual basis. This type of malware is reasonably easy to detect and eliminate, since its signatures and attack patterns are widely known and incorporated into anti-malware software. Even as it mutates in the wild (as *polymorphic malware* does), most anti-malware companies quickly notice these variations and add those signatures to their security suites.

Commodity malware is fairly common, unlike advanced malware that may be the product of advanced criminals or even nation-states. This type of malware specifically targets complex vulnerabilities or those that don't yet have mitigations, such as zero-day vulnerabilities, or advanced defenses. As such, advanced malware can be very difficult to detect and contain.

Anti-malware uses some of the same methods of detection that other security services and functions use. These methods include the following:

- Signature- or pattern-based detection for common malware.
- Behavior analysis and heuristic detection. Even if an executable is not identified as a piece of malware based on its signature, how it behaves and interacts with the system, other applications, and data may demonstrate that it is malicious in nature. Anti-malware solutions that use behavior and heuristic analysis can often detect otherwise unknown malicious code.
- Reputation-based services, where the anti-malware software communicates with vendors and security sites to exchange data about the characteristics of code. Based on what others have seen the code do, a reputation score is assigned to the code, enabling the anti-malware software to classify it as "good" or "bad."

The most important thing to remember about anti-malware solutions is that they must be updated on a consistent and continual basis with the latest signatures and updates. If an anti-malware solution is not updated frequently, it will not be able to detect new malware signatures or patterns. Most anti-malware solutions in an enterprise network are centrally managed, so updating signatures is relatively easy for the entire organization. However, administrators who are responsible for standalone hosts that use individually installed and managed anti-malware solutions must be vigilant about maintaining automatic updates or manually updating the anti-malware signatures often.

Unknown and potentially malicious code that is not detected by anti-malware solutions is a good candidate for reverse engineering. Reverse engineering is part of malware analysis, which means that an analyst obtains a copy of the potentially malicious code and analyzes its characteristics. These include its processes, memory locations, registry entries, file and resource access, and other actions it performs. This analysis also looks closely at any network traffic the unknown executable generates. Based upon this analysis, a cybersecurity analyst experienced in both programming and malware analysis may be able to determine the nature of the code.

Sandboxing

A *sandbox* is a protected environment within which an administrator can execute unknown and potentially malicious software so that those potentially harmful applications do not affect the network. A sandbox can be a protected area of memory and disk space on a host, a virtual machine, an application container, or even a full physical host that is completely separated from the rest of the network. Sandboxes have also been known over the years as *detonation chambers*, where media containing unknown executables were inserted and executed to study their actions and effects.

While anti-malware applications may be very effective at detecting malicious executables, attackers are also equally clever in obfuscating the malicious nature of those executables, simply by what is known as *bit-flipping* or changing the signature of the malware. A sandbox helps determine whether or not the application is malicious or harmless by allowing it to execute in

a protected environment that cannot affect other hosts, applications, or the network. Note that some anti-malware applications can automatically sandbox unknown or suspicious executables as part of their ordinary actions.

Machine Learning and Artificial Intelligence

Machine learning (ML) and artificial intelligence (AI) are advanced disciplines of computer science. While not specifically focused on cybersecurity, ML and AI tools, concepts, and processes can assist in analyzing very large heterogenous datasets. These technologies give analysts far more capabilities beyond simple pattern searching or correlation. These capabilities combine behavior analysis with complex mathematical algorithms and actually "learn" from the data ML- or AI-enabled systems ingest and the analysis those systems perform.

ML and AI, when integrated into a system such as a security orchestration, automation, and response (SOAR) implementation or a SIEM system, can be very helpful for looking at large volumes of data produced from sources all over the infrastructure. They can help determine if there are potentially malicious activities occurring, by enabling data correlation between seemingly unrelated data points, as well as keying in on obscure pieces of data that may be indicators of an otherwise difficult-to-detect compromise. This type of technology is very useful in threat hunting and looking for advanced persistent threat (ATP) presence in the infrastructure.

Both ML and AI can also be used for historical data analysis to determine exactly what occurred with a given set of data over a period of time. In addition, they can be used as predictive methods to determine future trends or potential threats.

Cross-Reference

Security orchestration, automation, and response (SOAR) is discussed at length in Objective 8.2, and security information and event management (SIEM) systems were discussed in Objective 7.2.

REVIEW

Objective 7.7: Operate and maintain detective and preventative measures In this objective we looked at various detective and preventive measures used in security operations. Most of these are technical controls designed to help detect anomalous or malicious activities in the network and prevent those activities from seriously impacting the organization. Preventive controls are critical in halting malicious activities before they even begin, but if preventive controls are not effective, detection is critical in stopping a malicious event.

Allow-listing and deny-listing are techniques used to permit or block network traffic, content, and access to resources based on rules contained in lists or rule sets. Items in an allow list are explicitly allowed, and any items not in the list are implicitly, or by default, denied. Items contained in a deny list are explicitly denied, and any items not in the list are implicitly, or by default, allowed. Most modern lists, however, contain both allow and deny rules.

Firewalls are traffic-filtering devices that can use various criteria (such as port, protocol and service) and deep content inspection to make decisions on whether to allow or deny traffic into, out of, or between networks. Network-based firewalls focus on network traffic, whereas host-based firewalls primarily focus on protecting individual hosts. Firewall types include packet-filtering, circuit-level, stateful inspection, and advanced next-generation firewalls. Newer firewall types include web application firewalls, whose purpose is to protect web application servers from specific attacks, and cloud-based firewalls, which function as a service offering from cloud providers and are more effective when most of the organization's infrastructure has been relocated to the cloud provider's data center.

Intrusion detection/prevention systems can detect and prevent attacks against an entire network (NIDS/NIPS) or individual hosts (HIDS/NIPS). IDS/IPSs can detect traffic using a number of methods, including signature- or pattern-based detection, behavior- or anomaly-based detection, and heuristic-based detection.

Third-party security services, also known as managed security services and Security as a Service, are often contracted to perform security functions that the organization is not staffed or qualified to perform. These services may include security device configuration, maintenance, and monitoring, log review, and SOC services. As long as the third-party service provider is trusted and a strong SLA is in place between the organization and the provider, this may be a preferred way of sharing risk. However, the risks of a third-party security service provider include unclear responsibilities, lack of reliability, undefined data ownership, and legal liability in the event of a breach.

A honeypot is a decoy host set up on a network to attract the attention of an attacker so that their actions can be recorded and studied, as well as to distract them from sensitive targets. A honeynet is a network of honeypot hosts.

Anti-malware applications can be deployed across the network or on individual hosts and are usually centrally managed. Anti-malware applications can detect malicious code using some of the same methods used for intrusion detection, such as signatures or patterns, changes in behavior, or even heuristic detection methods. Anti-malware can also use reputation-based scoring to determine if an unknown application may be a piece of malware. The most critical thing to remember about anti-malware solutions is that they are constantly being updated by vendors, so administrators must ensure that either automatic or manual updates occur on a frequent basis.

Sandboxing is a method of executing potentially unknown or malicious executables in a protected environment that is isolated from the rest of the network. This helps to determine whether the software is malicious or harmless without the potential of danger or damage to the infrastructure. Sandboxes can be virtual or physical machines.

Finally, we also examined the benefits of machine learning and artificial intelligence, which can allow analysts a much wider and deeper capability of analyzing massive amounts of disparate data to determine relationships and patterns.

7.7 QUESTIONS

1. You are a cybersecurity analyst in your company and are tasked with configuring a security device's rule set. You are instructed to take a strong approach to filtering, so you want to disallow almost all traffic that comes through the security device, except for a few select protocols. Which of the following best describes the approach you are taking?
   - **A.** Default allow
   - **B.** Default deny
   - **C.** Implicit allow
   - **D.** Explicit deny
2. You must deploy a new firewall to protect an online Internet-based resource that users access using their browsers. You want to protect this resource from injection attacks and cross-site scripting. Which of the following is the best type of firewall to implement to meet your requirements?
   - **A.** Packet-filtering firewall
   - **B.** Circuit-level firewall
   - **C.** Web application firewall
   - **D.** Host-based firewall

7.7 ANSWERS

1. **B** If you are only allowing a few select protocols and denying everything else, that is a condition where, by default, everything else is denied. The protocols that are allowed are explicitly listed in the rule set.
2. **C** A web application firewall is specifically designed to protect Internet-based web application servers, and can prevent various web-based attacks, including injection and cross-site scripting attacks.



Objective 7.8 Implement and support patch and vulnerability management

This objective addresses the necessity to update and patch systems and manage their vulnerabilities. This objective is also closely related to the configuration management (CM) discussion in Objective 7.3, as patches and configuration changes required to address vulnerabilities must be carefully controlled through the CM process.

Patch and Vulnerability Management

We discussed the necessity for secure, standardized baseline system configurations in Objective 7.3. Barring any intentional changes that we make to systems, they should, in theory, remain in those baselines indefinitely. However, this isn't possible in practice, since vulnerabilities to operating systems, applications, and even hardware are discovered on a weekly basis. Therefore, we must apply patches or updates to the systems, which means their baselines must change, sometimes frequently. Patch management can be a race against malicious actors who exploit vulnerabilities as soon as they are released. In some cases, vulnerabilities are exploited before there is a patch; these are called zero-day vulnerabilities and are among the most urgent vulnerabilities to mitigate.

Every organization must have a comprehensive patch and vulnerability management program in place. Obviously, policy is where it all begins. An organization's policies must address vulnerability management based on system criticality and vulnerability severity, include a patch management schedule, and address how to test patches before applying to hosts to mitigate any unexpected issues that may occur on production systems.

Managing Vulnerabilities

As mentioned, vulnerability bulletins are released on a weekly and even sometimes daily basis. Although the term "vulnerabilities" typically brings to mind technical vulnerabilities, such as those associated with operating systems, applications, encryption algorithms, source code, and so on, there are also nontechnical vulnerabilities to consider. Each vulnerability requires its own method of determining vulnerability severity and subsequent mitigation strategy.

Technical Vulnerabilities

Technical vulnerabilities are a frequent topic in this book. They apply to systems in general, but specifically can show up in operating systems, applications, code, network protocols, and even hardware, regardless of whether it is traditional IT devices or specialized devices such as those in the realm of IoT. Most technical vulnerabilities fall into one of a few categories, including, but not limited to

- Authentication vulnerabilities
- Encryption or cryptographic vulnerabilities
- Software code vulnerabilities
- Resource access and contention vulnerabilities

Technical vulnerabilities are often remediated with patches, updates, or configuration changes to the system, and in some cases completely new software or hardware is necessary to mitigate vulnerabilities if they cannot otherwise be eliminated. Technical vulnerabilities are usually detected during vulnerability scanning, which involves using a host- or network-based scanner to search a system for known vulnerabilities. We will discuss vulnerability scanning later in this objective.

Nontechnical Vulnerabilities

Nontechnical vulnerabilities can be more difficult to detect, and even harder to mitigate, than technical vulnerabilities, but they are equally serious. Nontechnical vulnerabilities include weaknesses that are inherent to administrative controls, such as policies and procedures, and physical controls. A policy addressing the mandated use of encryption is a serious weakness, for example, if no one is required to encrypt sensitive data. Physical vulnerabilities, such as lack of fencing, alarms, guards, and so on, can create serious security and safety concerns in terms of protecting facilities, people, and equipment.

Nontechnical vulnerabilities are also discovered during a vulnerability assessment, but this type of assessment looks more closely at processes and procedures, as well as administrative and physical controls. These vulnerabilities can't, however, be addressed by simply patching; more often than not, mitigating these vulnerabilities requires more resources, additional personnel training, or additional policies.

Managing Patches and Updates

Along with vulnerability management comes patching and update management. Technical vulnerabilities are most often addressed with patches or updates to operating systems and applications. However, patches and updates shouldn't simply be applied sporadically or only when an administrator has time. They should be carefully considered for both positive and potentially negative security and functional issues they may introduce into the infrastructure. Managing patches and updates includes considering the criticality of both the patch and the system, a solid patch update schedule, and formal patch testing and configuration management, all discussed next.



NOTE Although some professionals tend to use the terms "patches" and "updates" interchangeably in ordinary conversation, a patch is specifically used to mitigate a single vulnerability or fix a specific functional or performance problem. An update is a group of patches, released by a vendor on a less frequent basis (often scheduled periodically) and may add functionality to a system, or "roll up" several patches.

Patch and System Criticality

Criticality is a key concern in patch and update management, from two different perspectives:

- **System criticality** When installing patches and updates, critical assets, such as servers and networking equipment, may be offline for an indeterminate amount of time while the patch or update is applied and tested. Often this downtime is minimal, but for critical assets, the installation should be scheduled to meet the needs of the user base and the organization.

- **Patch and update criticality** The critical nature of the patch or update itself may be a factor. The patch may, for example, mitigate a zero-day vulnerability that creates high risk in the organization. The patch must be applied as soon as possible but should be balanced with the criticality of the systems that it must be applied to.

An organization must be prepared to make decisions that require balancing system criticality with patch criticality; this is often a subject that has to be addressed quickly by the entire change management board, so an organization must plan appropriately.



EXAM TIP Criticality of both patches and systems must be balanced when making the determination to install patches that mitigate serious vulnerabilities, especially those that have not been tested or may take systems down for an unknown period of time. You must balance the need to maintain system uptime and availability with the risk of not implementing the patch quickly.

Patch Update Schedule

Criticality aside, patches should be applied as soon as possible but should be prioritized according to the urgency of the patch and the criticality of the assets. An organization should have a regular schedule and routine for deploying patches, which should include the ability to test patches in a development environment before applying them to production systems. Administrators should have the opportunity to review patches based on the system criticality and the risk that leaving an unpatched system introduces, so they can appropriately schedule the patch for installation.

Routine patches may be applied only once a week or once a month but should at least be scheduled so that there is minimal disruption to the user base. This is where policy comes into play—the patch and vulnerability management policy should indicate a schedule based on criticality of patches and systems. A typical schedule may require critical patches to be applied within one business day and routine patches to be applied within seven calendar days. In any event, a regular patching schedule is absolutely necessary to make sure that new vulnerabilities are mitigated.

Patch Testing and Configuration Management

Simply downloading a patch and applying it to a production system is not a safe practice; patches and updates sometimes have adverse effects on systems, including affecting their functionality and performance, or even opening up further security vulnerabilities. Sometimes patching one vulnerability can lead to another, so patches should be tested on development or test systems prior to being implemented in the production environment. Unfortunately, sometimes the patch schedule makes this difficult, particularly when the patch or system is critical. In the event of an urgent patch, an organization might decide to apply it directly to production systems, after thorough research and ensuring the patch or update can be rolled back quickly if needed.

As mentioned previously, patches and updates are closely related to configuration management; sometimes applying a major patch or update changes the security baseline significantly. Once the patch has been tested and approved, it is implemented on production systems. For future systems, the patch may need to be considered for the initial build and included in the system master images. This requires configuration management and documenting changes to the official standardized baseline.

Patches and updates should be documented to the greatest extent possible; sometimes this may not be practical in the event of many patches that may come all at once, but at least maintaining a list of patches or a snapshot of the system state before and after patching can help later documentation. This is important because if a patch or update changes system functionality or lowers the security level of a system, documentation can provide valuable information when researching the root cause of the issue and can support potential rollback.

Cross-Reference

Configuration management was discussed at length in Objective 7.3.

REVIEW

Objective 7.8: Implement and support patch and vulnerability management In this objective we looked at patch and vulnerability management, which are closely related to system configuration and change management. We discussed the necessity to apply patches and updates to systems and applications on a scheduled basis and to consider both patch and system criticality when devising a patch management strategy. We also emphasized the importance of testing patches before implementing them on production systems, since even patches can cause systems to be less functional or less secure. We also examined vulnerability management, including the necessity to scan for and mitigate technical vulnerabilities. Nontechnical vulnerabilities may be more difficult to detect and mitigate than technical vulnerabilities, but addressing them is equally important. A proactive vulnerability and patch management program is critical to the security of the infrastructure.

7.8 QUESTIONS

1. Your company has ten servers running an important database application, some of which are backups for the others, with a significant vulnerability in a line-of-business application that could lead to unauthorized data access. A patch has just been released for this vulnerability and must be applied as soon as possible. You test the patch on development servers, and there are no detrimental effects. Which of the following is the best course of action to take in implementing the patch on all production servers?
   - **A.** Install the patch on all production servers at once.
   - **B.** Install the patch on only some of the production servers, while maintaining uptime of the ones that serve as backups.

C. Install the patch only on one server at a time.

D. Do not install the patch on any of the critical systems until users do not need to be in the database.

2. You install a critical patch on several production servers without testing it. Over the next few hours, users report failures in the line-of-business applications that run on those servers. After investigating the problem, you determine that the patch is the cause of the issues. Which one of the following would be the best course of action to take to quickly restore full operational capability to the servers as well as patch the vulnerabilities?

A. Rebuild each of the servers from scratch with the patch already installed.

B. Reinstall the patch on all the systems until they start functioning properly.

C. Roll back the changes and accept the risk of the patch not being installed on the systems.

D. Roll back the changes, determine why the patch causes issues, make corrections to the configuration as needed, test the patch, and install it on only some of the production servers.

7.8 ANSWERS

1. **B** Because the patch is critical, it must be installed as soon as possible and on as many servers at a time as practical. Since some servers are backups, those servers can remain online while the other servers are patched, and then the process can be repeated with the backups. Installing the patch on all servers at once would take down the production capability for an indeterminate amount of time and is not necessary. Installing the patch on only one server at a time would increase the window of time that the vulnerability could be exploited.

2. **D** The changes must be rolled back so that the servers are restored to an operational state, then research must be performed to determine why the patch caused issues. Any configuration changes should be investigated to determine if the issues can be corrected, and only then should the patches be tested and then reinstalled. You should install them only on some of the production servers so that some processing capability is maintained. If issues are still present, then you should repeat the process until the problem is solved. The servers should not need to be rebuilt from scratch, as this can take too long and is no guarantee that it will fix the problem. Reinstalling the patch over and over until the systems start functioning properly is not realistic. The risk of an unpatched system may be unacceptable to the organization if the vulnerability is critical.



Objective 7.9 Understand and participate in change management processes

As discussed in Objective 7.3, configuration management is a subset of the overall change management process. In this objective we turn to the change management program itself and discuss the change management processes that an organization must implement both to effectively manage changes that occur in the infrastructure and to deal with the security ramifications of those changes.

Change Management

Change management is an overall management process. It refers to how an organization manages strategic and operational changes to the infrastructure. This is a formalized process, intentionally so, to prevent unauthorized changes that may inconvenience the organization, at best, or, at worst, cripple the entire organization. Change management encompasses how changes are introduced into the infrastructure, the testing and approval process, and how security is considered during those changes.

Examples of the types of change the organization should pay particular attention to include

- Software and hardware changes and upgrades
- Significant architecture or design changes to the network
- Security or risk issues requiring changes (e.g., new threats or vulnerabilities)

Cross-Reference

Change management for software development is also discussed in Objective 8.1.

Change Management Processes

As a formalized management program, change management must necessarily have documented processes, activities, tasks, and so on. The change management program begins with policy (as does every other management program in the organization) that specifies individual roles and responsibilities for managing change in the organization. Change management processes also consider change as a life cycle and address both functional and security ramifications in that life cycle.

Change Management Policy

As mentioned, the change management program starts with policy. The organization must develop and implement policy that formalizes the change management program in the organization by assigning roles and responsibilities, creating a change management cycle, and

defining change levels. Roles and responsibilities are usually established and delineated by creating a change control board, discussed in the next section.

Change Management Board

The group of individuals tasked with overseeing infrastructure changes is called the change management board (CMB), change control board (CCB), or change advisory board (CAB). The individual members of the board are appointed by the organization's senior management and come from management, IT, security, and various functional areas to ensure that changes are formally requested, tested, approved, and implemented with the concurrence of all organizational stakeholders. Each identified stakeholder group normally appoints a member and an alternate to the board to ensure that group's interests are represented.

The CMB/CCB/CAB is often created by a document known as a *charter*, which establishes the various roles and their duties. The charter may also dictate the process for submitting and approving changes, the process for voting on changes, and how the board generally operates. The board usually meets on a regular basis, as dictated by the charter, to discuss and vote on an approved agenda of changes. These changes may come from the strategic plan for the organization, from changes in the operating environment and risk posture, or from changes necessitated by the industry or regulations.

Change Management Life Cycle

An organization should define a formal change management life cycle in policy. This life cycle should meet the requirements of the organization based on several criteria, including its appetite and tolerance for risk, the business and IT strategies for the organization, and urgency of changes that must take place. A generic life cycle for change management process in the organization includes these steps:

1. *Identify the need for change.* This can stem from planned infrastructure changes, results of risk or security assessments, environmental and technology changes, and even industry or market changes.
2. *Request the change.* The change must be formally requested (and championed) by someone in the organization—whether it is a representative of IT, security, or a business or functional area—who submits a formal business case justifying the change.
3. *Testing approval.* The CCB votes to approve or disapprove testing the change based on the request justification. The proposed change is tested to see how it may affect the existing infrastructure, including interoperability, function, performance, and security.
4. *Implement the change.* Based on testing results, the CCB may vote to approve the change for final implementation or send the change back to the requester until certain conditions are met. If the change is approved for implementation, the new baseline is formally adopted.
5. *Post-change activities.* These are unique to the organization and involve documenting the change, monitoring the change, updating risk assessments or analysis, and rolling back the change if needed due to unforeseen issues.



CAUTION Understand that these are only generic change management steps; each organization will develop its own change management life cycle based on its unique needs.

All changes aren't considered equal; some changes are more critical than others and may require the full consideration of the change board. Other changes are less critical and the decision to implement them may be routine and delegated down to a few members of the board or even IT or security if the changes do not present significant risk to the organization. All of these options and decision trees must be determined by policy. In this regard, an organization should develop change levels that are prioritized for consideration and implementation. While each organization must develop its own change levels, generally these might be considered as follows:

- **Emergency or urgent changes** These are changes that must be made immediately to ensure the continued functionality and security of the system.
- **Critical changes** These are changes that must be made as soon as possible to prevent system or information damage or compromise.
- **Important changes** These are changes that must be performed as soon as practical but can be part of a planned change.
- **Routine changes** These are minor changes that could be made on a daily or monthly basis; most noncritical patching or updates generally fall into this category.



EXAM TIP The priority of changes is critical for an organization to develop and adhere to, since changes with some urgency must be quickly considered and implemented, especially those that impact the security of the infrastructure.

Note that some organizations, in addition to prioritizing changes based on urgency, also categorize changes in terms of the effort required to implement the change or the scope of the change. Examples may include categories such as major changes and minor changes.

Security Considerations in Change Management

Security is a critical factor in proper change management. Any changes to the infrastructure introduce a degree of risk, so these changes must undergo an assessment to determine the extent of security impact or risk they will introduce to the organization's assets. Both before and after a change is implemented, the organization should perform vulnerability testing on the affected systems to see what deltas in vulnerabilities the changes will produce; this way any new vulnerabilities can be properly attributed to the change itself rather than other factors.

The affected systems should be "frozen" or locked in configuration before this occurs, so as to not introduce any unexpected or random factors into the testing and subsequent changes. Once a security impact of the change is determined, this information is considered along with other factors, such as downtime, scope of the change, and urgency, before finalizing a change.

REVIEW

Objective 7.9: Understand and participate in change management processes This objective introduced the concept of change management as a formal program in the organization. Change management means the organization must have a formalized, documented program in place to effectively deal with strategic and operational changes to the infrastructure. Change management begins with a comprehensive policy that outlines roles, responsibilities, the change management life cycle, and categorization of changes. The change management board is responsible for overseeing the change process, including accepting change requests, approving them, and ensuring that the process follows standardized procedures. The change life cycle includes requesting the change, testing the change, approval, implementation, and documenting the change. Security impact considerations must be included in the change management process since change may introduce new vulnerabilities into the infrastructure.

7.9 QUESTIONS

1. Which of the following formally creates the change management board and establishes the change management procedures?
   - **A.** Security impact assessment
   - **B.** Charter
   - **C.** Policy
   - **D.** Change request
2. Your company's change management board is evaluating a request to add a new line-of-business application to the network. Which step of the change management life cycle should be performed before final approval of this request?
   - **A.** Document the application and its supporting systems.
   - **B.** Test the changes to the infrastructure in a development environment.
   - **C.** Perform a rollback to the original infrastructure configuration.
   - **D.** Submit a formal business case to the board from the responsible business area.

7.9 ANSWERS

1. **B** A change management board charter is often used as the source document to create the change board and establish its processes.
2. **B** Before final approval of the change to the infrastructure, all changes should be tested in a development environment.



Objective 7.10 Implement recovery strategies

In this objective we begin our discussion of disaster recovery and business continuity. We will discuss various recovery strategies, including those associated with backups, recovery sites, resilience, and high availability. Although the recovery strategies we will cover are often associated with disasters in particular, these same strategies can also be used during a variety of incidents, as we discussed in Objective 7.6. Because of this, Objective 7.10 serves as an important link between our previous discussion on incident management and Objective 7.11 that addresses disaster recovery planning, which we will discuss later on during this domain.

Recovery Strategies

Recovery strategies are designed to keep the business up and functioning during a disaster or incident and to expedite a return to normal operations. The key to recovery strategies are resiliency, redundancy, high availability, and fault tolerance. We will discuss each of these in this objective, as the different strategies that can be implemented by an organization usually target one of these key areas.

Backup Storage Strategies

To recover an organization's processing capabilities, the organization must ensure that its data is backed up prior to any adverse event that requires data restoration. There are several backup storage strategies an organization can use; the decision regarding which one (or combination of several) to use depends on a few factors, including

- How much data the organization can afford to lose in the event of a disaster
- How much data the organization requires to restore its processing capability to an acceptable level
- How fast the organization requires the data to be restored
- How much the backup method or system and media cost
- How efficient the organization's network and Internet connections are in terms of speed and bandwidth

In the next few sections we will discuss various backup strategies that vary in cost, speed, recoverability, and efficiency.

Traditional Backup Strategies

Traditional backup strategies involve backing up individual servers to a tape array or even to separate disk arrays. Because these backup methods are often expensive and time consuming, a system of using backup strategies such as full, incremental, and differential backups was developed to deal with these issues. Briefly, the strategies are executed as follows:

- **Full backup** The entire hard disk of a system is backed up, including its operating systems, applications, and data files. This type of backup typically takes a much longer amount of time and requires a great deal of storage space.
- **Incremental backup** This strategy backs up only the amount of data that has changed since the last full backup. It requires less storage space and is somewhat faster, but the data restore time can take much longer. The last full backup and each incremental backup since the last full backup have to be restored one at a time, in the order in which they occurred. Incremental backups reset the archive bit on a backup file, showing that the data has been backed up. If the archive bit is set to "on," as happens when a file changes, this means it has not been backed up.
- **Differential backup** This backup strategy involves backing up only data that has changed since the last full or incremental backup; the difference between this type of backup and an incremental backup is that the archive bit for the backup files is *not* turned off. In the event of a restore, only the last full backup and the last differential backup have to be restored, since each subsequent differential backup includes all the previous changed data. However, as differential backups are run, they became larger and larger, since they include all data that has changed since the last full backup.

Again, these strategies were devised during the days when backup solutions were highly expensive, backup media was slow and unreliable, and backups were very time-consuming and tedious. Although there are still some valid uses of these strategies today, they have largely become obsolete due to the much lower cost of other backup media, high-speed networking, and both the availability and inexpensive nature of technologies such as cloud storage, all of which are discussed in the following sections.

Direct-Attached Storage

Direct-attached storage is the simplest form of backup. It also may be the least dependable, since it is an external storage media, such as a large-capacity USB drive, directly attached to the computing system. While direct-attached storage may be fast enough for the organization's current requirements, it is unreliable in that a physical disaster that damages a system may also damage the attached storage. This type of storage is also susceptible to accidents, intentional damage, and theft. Direct-attached storage should never be used as the only form of backup for critical or sensitive data, but it may be effective as a secondary means of backup for individual user workstations or small datasets.

Network-Attached Storage

A network-attached storage (NAS) system is the next step up from direct-attached storage. It is simply a network-enabled storage device that is accessible to network hosts. It may be managed by a dedicated backup server running enterprise-level backup software. It may also double as a file server. Over a high-speed network, NAS can be quite efficient, but suffers from some of the same reliability issues as direct-attached storage. Any disaster that occurs causing damage to the infrastructure may also damage the NAS system. NAS may be sufficient for small to medium-sized businesses and may or may not offer any redundancy capabilities.

Storage Area Network

A storage area network (SAN) is larger and more robust than a simple network-attached storage device. A SAN is designed to be a significant component of a mature data center supporting large organizations. It may have its own cluster of management servers and even security devices to protect it. A SAN is often built for redundancy by having multiple storage devices that fail over to each other. Components are connected to each other and to the rest of the network with high-speed fiber connections. While designed for high availability and efficiency, a SAN still suffers from the possibility that a catastrophic event damaging the facility can also damage the SAN.

Cloud Storage

Cloud-based storage is becoming more of the norm than on-premises storage. While on-premises storage is still necessary for short-term or smaller-level data recovery, the remote storage capabilities of the cloud support large-scale data recovery that is fast and reliable. Even if a facility is entirely destroyed, data that is backed up over high-speed Internet connections to the cloud is easily available for restoration. The only limiting factors an organization may face when using cloud-based storage as a disaster recovery solution are the associated costs of cloud storage space, which may increase if the amount of data increases over time, as well as the availability of high-speed bandwidth from the organization to the cloud provider.

Offline Storage

Offline storage simply means that data is backed up and stored off of the network and/or at a remote location away from the physical facility. The methods for creating and using offline storage may be manual or electronic; even cloud-based storage is considered offline since it is not part of the same organizational infrastructure and is not housed in the same physical facility. Traditionally, organizations manually transported backup media, such as tapes, optical discs, and hard drive arrays, to a geographically separated site so that if a disaster damaged or destroyed the primary processing facility, the backup media would still be available. Another benefit of offline storage is that a malware infection, ransomware, or other malicious attack that impacts the primary processing facility is far less likely to impact the offline storage.

Electronic Vaulting and Remote Journaling

Contrary to traditional thinking, backups don't always have to be on a full, incremental, or differential basis. Backups can also be performed on individual files and even individual transactions. They can also be captured electronically either on a very frequent basis or in real time. *Electronic vaulting* is a backup method where entire data files are transmitted in batch mode (possibly several at a time). This may happen two or three times a day or only during off-peak hours, for instance. Contrast this to *remote journaling*, which only sends changes to files, sometimes in the form of individual transactions, in either near or actual real time.

Both methods require a stable, consistent, and sometimes high-bandwidth network or Internet connection, depending on if the backup mechanism is local within the network or located at a geographically separated site.



EXAM TIP Keep in mind the differences between electronic vaulting and remote journaling. Electronic vaulting occurs on a batch basis and moves entire files on a frequent basis. Remote journaling is a real-time process that transmits only changes in files to the backup method or facility used.

Recovery Site Strategies

A recovery site, sometimes referred to as an alternate processing site, is used to bring an organization's operations back online after an incident or disaster. A recovery site is usually needed if an organization's primary data processing facility or work centers are destroyed or are otherwise unavailable after a disaster due to damage, lack of utilities, and so on. Many considerations go into selecting and standing up a recovery site. First, an organization must evaluate how much resources, such as money and equipment, to invest based on the likelihood that it will need to use an alternate recovery site. For example, if your organization has performed risk assessments and adequately considered the likelihood and impact of catastrophic events, and determined that the odds of such an event are very low, it very well may decide that maintaining a recovery site that is staffed 24/7 and has fully redundant equipment is unnecessary and a waste of money. Another consideration is how long your organization can afford to be down or operating at minimal capacity. If it can survive a week without being operational, then having a recovery site available that can be activated over a few days likely will be more cost-effective than having one that's ready to go at a moment's notice. In regard to both of these considerations, recovery site strategies depend very much on how well your organization has performed a risk analysis.

Cross-Reference

Risk analysis was covered in depth in Objective 1.10.

Another consideration in creating a strategy for a recovery site is the ease with which the organization can activate the recovery site and relocate its operations there. If the site is a considerable distance away, relocation might be too difficult if a natural disaster disrupts roads,

transportation, employee family situations, and so on. Those are the types of issues that might prevent an organization from efficiently relocating to another site and adequately recovering its operations.

Other issues that your organization should consider when formulating a recovery site strategy include

- Potential level of damage or destruction caused by a catastrophic event
- Possible size of the area affected by a disaster
- Nature of the disaster (e.g., tornado, flood, hurricane, terrorist attack, war, etc.)
- Availability of public infrastructure (e.g., widespread power outages, damage to highways and transportation modes, hospital overcrowding, food shortages, etc.)
- Government-imposed requirements, such as martial law, protections against price gouging, quarantine zones, etc.
- Catastrophes that negatively affect employee families, since this will definitely impact the workforce
- Likelihood of multiple organizations competing for the same or similar resources

In addition to these considerations, recovery sites must have sufficient resources to support the organization; this includes utilities such as power, water, heat, and communications. There must be enough space in the facility to house the employees needed for recovery actions, as well as the equipment that may be relocated. These are constraints an organization will need to take into account when selecting recovery sites, discussed next.

Multiple Processing Sites

An organization must carefully consider several questions when selecting the right type of alternate processing site for its needs:

- How soon after operations are interrupted will the organization need to access the site?
- Will the site have to be fully equipped and have all the proper utilities within a few hours, or can the business afford to wait a few days or weeks before relocating?
- How much time, money, and other resources can the organization afford to invest in an alternate processing site before the cost outweighs the risk that the site will be needed?

These are all questions that can be answered with careful disaster recovery and business continuity planning.

Traditional alternate processing sites include cold, warm, and hot sites; these are used when an organization must physically relocate its operations due to damage to its facilities and physical equipment. This damage may mean that the facility is unusable because it lacks structure, lacks utilities, or even presents safety issues. The other types of sites we will examine, including reciprocal, cloud-based, and mobile sites, don't necessarily require an organization to relocate its physical presence, instead providing "virtual" relocations or simply alternate processing capabilities.

Cold Site

A cold site is not much more than empty space in a facility. It doesn't house any equipment, data, or creature comforts for employees. It may have very limited utilities turned on, such as power, water, and heat. It likely will not have high-bandwidth Internet connections or phone systems the organization can immediately use. This type of site is used when an organization either has the luxury of time before it is required to relocate its physical presence or simply cannot afford anything else. As the least expensive alternate processing site option, a cold site is not ready to go at a moment's notice and must be furnished, staffed, and configured after the disaster has already taken place. All of its equipment and furniture will have to be moved in before the site is ready to take over processing operations.

Warm Site

A warm site is further along the spectrum of both expense and recoverability speed. It is more expensive than a cold site but can be activated sooner in the event of a disaster to restart an organization's processing operations. In addition to space for employees and equipment, a warm site may have additional utilities, such as Internet and phone access, already turned on. There may be a limited amount of furniture and processing equipment in place—typically spares or redundant equipment that has already been staged at the facility. A warm site usually requires systems to be turned on, patched, securely configured, and have current data restored in order for them to be effective in taking over operations.

Hot Site

A hot site is the most expensive option for traditional physical alternate processing sites. A hot site is ready for transition to become the primary operating site quickly, often within a matter of minutes or hours. It already has all of the utilities needed, such as power, heat, water, Internet access, and so on. It usually has all of the equipment needed to resume the organization's critical business processes at a moment's notice. In a hot site scenario, the organization also likely transfers data quickly or even in real time to the alternate site without the risk of losing any data, especially if it uses high-speed data backup connections. Because data can be transferred in large volumes quickly with a high-speed Internet connection, many organizations use their hot site as their off-site data backup solution, which makes spending the money on a physical hot site much more efficient and cost-effective.



EXAM TIP Traditional alternate processing sites should be used when the organization needs to physically relocate its operations. A cold site is least expensive but requires the most time to be operational; a warm site is more expensive but can be ready faster; and a hot site is the most expensive and can be ready to take over processing operations within hours. The decision regarding which of these three types of sites to use is based on how much the organization can afford and how fast it needs to be operational.

Reciprocal Sites

Many organizations have a reciprocal site agreement with another organization that specifies each organization can share the other's resources in the event of a disaster that affects only one of the organizations. This type of agreement gives the affected organization an opportunity to recover its operations without having to move to a traditional cold, warm, or hot recovery site.

A reciprocal site agreement may be an effective strategy, but there are a few key considerations:

- Are the organizations direct competitors, in a similar business market, or offering similar products or services? On one hand, being in a similar market means the two organizations probably serve well as reciprocal sites since each organization likely uses comparable equipment, applications, and processing capabilities. On the other hand, if they are direct competitors, it might not be a suitable long-term solution for maintaining confidentiality of proprietary or sensitive information.
- Are the organizations in the same geographic area? If so and a natural disaster strikes, neither organization may be able to support the other.

Cloud Sites

Cloud service providers offer a new opportunity for organizational resilience. Traditionally, if an organization suffered physical facility damage, it had to find a new place to set up operations. Depending on the level of damage to the operations, recovery could require long hours provisioning new servers, restoring data from backups (still possibly losing a couple of days' worth of data in the process), reloading applications, and so on. Cloud computing changes this entire paradigm. If an organization suffers serious physical harm to its facilities or equipment, it's possible to have complete redundancy for its systems and data built into the cloud. Cloud-based redundancy means that the only reasons an organization would have to find an alternate processing location would be to preserve the health, safety, and comfort of its personnel. With the surge in remote working necessary to maintain operations during the global COVID-19 pandemic, use of cloud solutions has greatly accelerated.

Even during normal processing times, with no risk of disaster or catastrophe, organizations were already slowly moving a great deal of their processing power to the cloud. Organizations have increasingly moved to Software as a Service (SaaS), Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and other, "Anything as a Service (XaaS)" cloud offerings. So if a disaster strikes, the organization may not even have to expend much effort toward disaster recovery or business continuity activities. If the majority of your organization's systems are already functioning in the cloud, the only ones that need to be recovered are likely lower-priority systems or legacy systems that have not transitioned to the cloud.

Cross-Reference

Cloud-based systems were covered in depth in Objective 3.5.

Mobile Sites

Mobile sites add another dimension to alternate processing capabilities. While cloud-based services can certainly support information processing operations, an organization still may require a physical alternate processing site. If an organization determines, for example, that a cold or warm site is insufficient for bringing operations online quickly enough and that a hot site is too expensive, a mobile site may offer a convenient, economical alternative. A mobile site can be built into a large van, a bus, or even an 18-wheeled transfer truck. While this type of alternate site won't hold many employees, it gives organizational leadership and key personnel the ability to still work together from a physical "command post." The mobility advantages are clear; the mobile site can travel away from the major disaster area, where there may be plentiful power, fuel, and other resources, and where the infrastructure may be more supportive for recovery. Larger organizations may own their own vehicle specially outfitted as a mobile site; smaller businesses may be required to lease such a specialized vehicle.

Mobile sites are considered miniature hot sites; while a mobile site may not have the capacity of a large building or other facility, it can certainly hold enough physical equipment to maintain a small data center, particularly if many of its hosts are virtualized or actually present in the cloud and accessible through a strong Internet connection.

Resiliency

Resiliency is the capability of a system (or even an entire organization) to continue to function even after a catastrophic event, although it may only function at a somewhat degraded level. For example, a system with 32GB of RAM that suffers an electrical problem and loses half of that RAM is still functional, although it may be limited in its processing capability and run slower than normal. The same can be said of a server that has dual power supplies and loses one of them or experiences a failure of a disk in a hardware raid array. These components may not necessarily fail, but they may operate at a marginal level of capability.

In the case of an entire organization, resiliency means it may lose some of its overall capabilities (people, equipment, facilities, etc.) but still be able to function at an acceptable, albeit degraded, level. High resiliency is one of the goals of business continuity; it is enabled by having redundant (and often duplicate) system components, alternative means of processing, and fault-tolerant systems.

High Availability

An organization that has implemented high availability (HA) can expect its infrastructure and services to be available on a near constant basis. Traditionally, high availability means that any downtime experienced by the processing capability or a component in the infrastructure is limited to only a few hours or a few days a year. In the early days of e-commerce, most businesses could afford that level of downtime. Only critical services required higher availability rates. For example, if the infrastructure had an availability rate of 99.999 percent (commonly referred to as "five nines uptime"), it would only be down, theoretically, just five minutes and 15 seconds per year. Typically, only very large organizations with critical services could afford this level of availability.

With the advent of high-speed Internet, cloud technologies, virtualization, and other technologies, high availability is a far less expensive prospect, affordable by even small businesses. Additionally, in the ultra-connected global age we live in, even a five-minute period of downtime per year may be too much. Consider the millions of transactions that may occur in a single minute that could be lost with even a small amount of downtime. Fortunately, the technologies available for resiliency and redundancy almost guarantee, if properly implemented, that even a small level of downtime can be almost eliminated.

Quality of Service

Quality of service (QoS) is a somewhat subjective term. Essentially, QoS is the minimum level of service performance that an organization's applications and systems require. For example, the organization could establish a minimum level of bandwidth required to move its data in and out of, as well as within, the organizational infrastructure. Different types of data and the context in which they move throughout the organization affect the level of service quality required. For example, high-resolution video typically requires a higher level of bandwidth than simple text; any degradation of bandwidth or network speeds reduces the quality of the video or prevents it from being sent or viewed.

In a situation such as a disaster that results in the loss of bandwidth, the organization may have to accept that it cannot move those large, high-resolution video files across the network and may have to settle for smaller files with lower resolution that still get the job done. For instance, Voice over IP (VoIP) traffic usually gets top priority in a network, as it can't tolerate low bandwidth that causes interruptions (called jitter), unlike e-mail services. Users may not notice a 50ms+ delay in e-mail traffic, but voice and video traffic will experience jitter and quality degradation, so there is a minimum bandwidth needed for those services and applications. So, often QoS is a determination of what the minimum service levels the organization *needs* for different services versus what it normally has. QoS is also improved by redundant or alternative capabilities, fault tolerance, and service availability.

Fault Tolerance

Fault tolerance means that the infrastructure or one of its systems is resistant to failure. The expectation is that if a network has high fault tolerance, it can resist complete failure of one or more components and still function. The following are a few ways to assure fault tolerance:

- Invest in higher-quality components that have lower failure rates. Cheaply made components often break more often even under lighter loads.
- Invest in redundant components, such as servers with dual power supplies, mirrored RAID arrays, or multiple processors. A stronger example is server clustering using virtual machine capabilities.

The decision to invest in fault-tolerant components and designs can help ensure higher overall availability.

REVIEW

Objective 7.10: Implement recovery strategies In this objective we discussed various recovery strategies. All of these strategies address key concepts such as resiliency, redundancy, high availability, fault tolerance, and quality of service.

Backup storage strategies are chosen based on a number of factors, including cost, speed at which the organization needs to restore data, bandwidth and speed available for network and Internet connections, how much data the organization needs to restore and how quickly it must be restored, as well as how much data the organization can afford to lose during a disaster.

- Backup solutions include traditional backup methods that use tape or hard disk arrays and are performed using full, incremental, or differential strategies.
- Direct-attached storage is a device physically attached to a system, such as a USB hard disk.
- Network-attached storage is a storage appliance connected to the network and accessible by various systems.
- Storage area networks (SANs) are typically larger, more robust storage arrays consisting of multiple devices and connected by a high-speed backbone.
- Direct-attached storage, network-attached storage, and SANs all have the vulnerability that if the entire facility is damaged or destroyed, they will also be affected.
- Cloud-based storage is not impacted by damage to an organization's facility, although it may be temporarily inaccessible due to network outages in the facility. However, cloud-based storage provides almost a perfect backup solution if the organization does not have to physically relocate.
- Offline storage means that data is stored at a remote site or in the cloud, using manual or electronic means.
- Electronic vaulting involves batch processing of entire files on a frequent basis.
- Remote journaling is performed in real time and only requires piecemeal backups of files, usually by transactions.
- Recovery site strategies center on how much an organization can afford, its risk of needing an alternate processing site, and how fast it needs to reach operational status after a disaster.
- A cold site is essentially empty space with minimal utilities; it does not offer the capability to recover quickly after a disaster, but it is the least expensive alternate processing site option.
- Warm sites offer a midway point along the spectrum of expense and recoverability; they are more expensive than cold sites, but include additional utilities, some equipment on standby, and the ability to get an organization up and running somewhat faster than a cold site.

- A hot site is the most expensive type of physical alternate processing site since it offers a fully functional physical processing space with redundant equipment and data, enabling an organization to return to operational capacity within minutes or hours.
- Cloud sites are a great option for organizations that have already migrated some of their processing capability to cloud-based services; they offer a fairly complete recovery solution for organizations that do not need to physically relocate to an alternate processing site.
- An organization can use a mobile site if it needs to maintain a minimal physical presence after a disaster but requires the alternative space more quickly than a cold or warm site affords and more cheaply than a hot site costs. The mobile nature allows the organization to move its command post or base of operations out of the danger zone or other disaster area to an area where there is better utility and infrastructure support. However, a mobile site does not supply adequate space for large numbers of people.
- Recovery strategies center on key concepts such as resiliency, high availability, quality of service, and fault tolerance.
- Resiliency means that a component, system, or the entire infrastructure of the organization will not fail completely or easily; the processing capability may be reduced to a lower level but will still be functional.
- High availability means that systems and data must be available on a near constant basis. With today's need for massive amounts of data processed almost in real time, even downtime of a few seconds can be catastrophic for an organization. Fortunately, modern technology such as cloud services, high-speed networking, virtualization, and quality components can help assure high availability even for small businesses.
- Quality of service prioritizes bandwidth for selected systems or data.
- Fault tolerance is the resistance to failure by components in the infrastructure. Fault tolerance is made possible through equipment and component redundancy, duplicate capabilities, and the use of quality equipment.

7.10 QUESTIONS

1. Which of the following traditional backup methods only backs up data that has changed since the last full backup, which also resets the archive bit for the data?
   - **A.** Incremental
   - **B.** Full
   - **C.** Differential
   - **D.** Transactional

2. Which of the following traditional backup sites should be used for physically relocating an organization's processing capability and personnel in the fastest manner possible, with all needed equipment and data already prepositioned?
   - **A.** Cold site
   - **B.** Warm site
   - **C.** Hot site
   - **D.** Mobile site

7.10 ANSWERS

1. **A** An incremental strategy backs up data that has changed since the last full backup only. After the data is backed up, it resets the archive bit, showing that data has been backed up. During a data restore situation, first the full backup is restored, and then each sequential incremental backup must be restored.
2. **C** A hot site is the most appropriate type of alternate processing site for this scenario, since the organization must be up and running quickly, and the site must contain all necessary equipment and data needed to restore full operations.



Objective 7.11 Implement Disaster Recovery (DR) processes

Objective 7.11 continues our discussion of how organizations must plan for and react to disasters. Disaster recovery is focused on saving lives first, then equipment, facilities, and data. The recovery strategies discussed in Objective 7.10 not only apply to general incidents and business interruptions, but also to disasters that may destroy facilities, damage equipment, place personnel in harm's way, and seriously disrupt business operations. This objective discusses the processes that go into a disaster recovery plan and how they are developed.

Disaster Recovery

As mentioned previously in this domain, there is a blurry line between incident response (IR), disaster recovery (DR), and business continuity (BC) activities. There are similarities and commonalities between all three of these major activities, such as planning and response; in some cases you will be performing the same type of activity or task for any of the three areas. Context is sometimes the only differentiating factor between these areas. The major differentiator for disaster recovery is that it is focused foremost on saving lives and preventing harm to individuals, and then recovering or salvaging equipment, systems, facilities, and even data. In this objective we will look at the various processes that go into planning and executing the disaster recovery plan.



EXAM TIP The primary focus of disaster recovery is saving lives and preventing further harm to individuals. The secondary focus is saving, recovering, or salvaging equipment, systems, information, and facilities. Getting the business back up and running is the focus of business continuity, not disaster recovery; however, many of their processes overlap and are executed simultaneously as the organization is able to do so.

Saving Lives and Preventing Harm to People

People are the most important focus of the disaster recovery plan. The primary goal of disaster recovery is saving lives and ensuring the safety of individuals. There are several key elements that go into ensuring safety and preventing harm.

- Establish solid emergency procedures, such as those related to evacuations, fire safety, tornado or earthquake sheltering, and, unfortunately necessary, manmade catastrophes, such as active shooters, terrorist attacks, and serious crime against the organization.
- Be prepared with the proper lifesaving and emergency safety equipment, including fire extinguishers, first-aid kits, stretchers, and so on. In the event of a serious disaster, if personnel have serious injuries and need significant medical attention, the initial response is even more critical. Options should be considered in the event of a mass casualty event to triage injuries and transport people to the closest medical facility immediately.
- Training personnel is also a critical factor in preserving lives. Personnel should be trained on emergency situations, as discussed later in this objective. Emergency procedures should be exercised; fire evacuation drills, active shooter exercises, and other disaster scenarios must be practiced, ensuring that people understand the plans and their responsibilities during these difficult situations.

The Disaster Recovery Plan

The success of the disaster response itself is highly dependent on a good disaster recovery plan (DRP). The planning process is designed to ensure that in the event of a disaster, human lives are saved and equipment is recovered or salvaged. It's important that the right amount of effort be put into developing the DRP, as planning also extends to ensuring that the right people, processes, equipment, supplies, and other resources are in place and functional.

The DRP must address several key points, including

- Criteria for declaring a disaster
- Preserving human lives and ensuring personnel safety

- Procedures for activating the disaster recovery team
- Communication strategies for all critical personnel
- Lines of authority and responsibilities of key personnel
- Procedures for assessing damage and recoverability of assets
- Processes for initiating business continuity activities
- Exercising contingency plans

Most of these topics have been discussed in the preceding objective or will be covered in this objective and the next two objectives, since all of these objectives relate to business continuity planning and disaster recovery planning.

Response

The DRP addresses response actions that the organization must be adequately prepared to take. This involves a wide gamut of actions, which include:

- Steps for declaring a disaster
- Activating the disaster response team
- Determining the safety and status of all personnel
- Establishing communications and command structure for the recovery efforts
- Damage assessment for all facilities and critical equipment
- Recovering or salvaging equipment, systems, and data
- Establishing alternate facilities to house personnel for business continuity efforts
- Restoring utilities such as power, water, heat, and communications
- Implementing business continuity efforts

Personnel

The organizational personnel selected to be part of the response team should be trained and qualified to perform all activities detailed in the DRP. Not only do these activities include general emergency procedures, but disaster recovery team personnel should also be trained in activities such as facilities and equipment damage assessment, recovery and salvage operations, and restoring business operations.

Communications

Communication is vital during disaster recovery efforts. When an organization creates its disaster recovery plan, it must define its communications procedures and strategy prominently in the plan. In addition to the obvious need to communicate up the chain to management and senior leaders about the recovery situation as it plays out, the DR team's leaders need to be able

to communicate with members of their team, other employees, and any other stakeholders. Communications also happen laterally across the organization. Different functional areas will likely need to communicate with each other during the recovery effort to coordinate activities.

Internal communications are not the only concern. Since a disaster likely involves many outside agencies and stakeholders, external communications are also critical. The communications plan should dictate primary and alternate communications personnel who will pass information on to external agencies. Depending on the size and complexity of the organization, there may be a specific person (or more than one person) designated to pass information to the media, customers, business partners, suppliers, law enforcement, regulatory agencies, emergency personnel, and so on. Senior leadership must also dictate how much information can be shared with specific external parties. In the event of a terrorist attack or a crime, for example, only certain information should be shared with the public, but all available information the organization has must be shared with law enforcement.

During the initial phases of disaster recovery, land lines, network communications (e.g., e-mail and instant messaging), or even cellular services may be out of commission. The organization may have very few options for communicating with its personnel to ensure their safety and to initiate recovery operations. Radios are excellent contingency communications methods. They must be purchased and tested (routinely) in advance of any disaster and issued to key personnel. Normally, radios used during contingencies will be of sufficient power and range to cover longer distances and operate on required emergency or general use frequencies. The organization may find that in the most serious of disasters, physically sending designated personnel, called *runners*, to personal residences to relay information may be necessary. Secondary failures of infrastructure, such as damage to highways, loss of power, and other concerns, may even interrupt sending personnel physically to communicate with others. These alternate methods of communicating with organizational personnel should be established in the DRP.

The organization should establish procedures such that, absent communications abilities after a disaster, personnel should shelter in place or attempt to communicate with their supervisors or key members of the disaster recovery team as practical. The organization may have to accept delays in communication if disaster conditions simply prohibit it, but it should have communications plans and procedures in place for when conditions permit contacting personnel and initiating disaster recovery efforts.



EXAM TIP An organization must consider several different aspects of communications when developing its DRP: communications up and down the management chain, communications laterally across the organization, and communications with external stakeholders.

Assessment

Assessment is a critical task that immediately precedes recovery and business continuity.

An assessment team, composed of members of the disaster recovery team, is responsible for determining the following, among other things:

- The safety of facilities for return of personnel
- The viability of equipment and systems for functionality
- If damage to equipment or systems can be repaired
- Whether any information left behind on storage media in the facility is intact
- If utilities and other critical resources are still functional or present

The assessment team is responsible for going to an organization's facilities, as soon as conditions permit, and determining how much of the organization's equipment, facilities, systems, and information are recoverable. While this must be done as soon as possible after a disaster, the team must wait until conditions are safe or stable. The assessment team should be composed of experts who have the knowledge and experience necessary to assess the situation and determine how much of the organization's assets, especially critical ones, can be recovered and restored.

Note that some determinations, such as the safety of facilities, may require outside expertise from municipal building inspectors, fire department safety personnel, and so on. For disasters that strike a larger area, local government officials may have the authority to decide whether or not an organization can begin recovery operations.

Restoration

Restoration occurs after disaster conditions have improved sufficiently to permit safe movement and activity. A prolonged or severe disaster may prevent restoration efforts indefinitely due to unsafe conditions, lack of infrastructure, and general chaos.

Restoration involves getting an organization from the point of being damaged and harmed by the disaster to a point where the organization is ready to begin business continuity efforts. The goal is to restore the environment to a condition in which personnel can safely resume working and accessing the resources they need. This may also include considering the personal lives of organizational personnel, since that will directly impact their availability and the effectiveness in helping to get the organization back on its feet.

Training and Awareness

The importance of training and awareness programs that teach people the critical tasks they must perform during the disaster recovery effort cannot be overemphasized. If people are not aware of even the fundamentals of disaster recovery activities, they definitely will not be able to perform those tasks in the chaos and uncertainty of an actual disaster.

Fundamental DR information that should be taught to personnel incorporates the nature of disasters, what's expected during the recovery effort, and how they can save lives, maintain safety, and then perform tasks to recover equipment, data, and facilities. Basic lifesaving and safety training includes CPR, first aid, emergency evacuation procedures, and so on. All personnel in the organization should receive this minimal level of training, even if they are not assigned to any disaster recovery responsibilities.

Employees who are assigned DR responsibilities should be trained as primary on some tasks and alternate on other tasks, since it is likely they may be needed to perform multiple functions during the course of a recovery effort. For example, the damage assessment team should be trained on how to safely go through a facility, once it has been cleared by the appropriate safety or law enforcement personnel, and inventory equipment, systems, and information and assess whether or not they are still viable. Those same team members, however, likely won't be performing that function until personal safety is assured, so they can be trained to perform other tasks that occur prior to damage assessment.

Classroom training must be a part of the training and awareness program for disaster recovery; however, exercising the disaster recovery plan is also of critical importance, as will be discussed in Objective 7.12. Simply teaching people what they need to do is insufficient—they must practice it in order to become proficient. It must become like muscle memory so that their responses are automatic and accurate. As we will discuss in Objective 7.12, exercising the DRP also helps to weed out inconsistencies, issues, and conflicts within the plan. The people performing tasks and activities to support the plan will be part of finding these issues, so they will be much more aware of them and how they impact their performance during the recovery effort.

Lessons Learned

As we discussed in Objective 7.6 in the context of incident management, lessons learned are critical in taking all available information learned during a negative event and using that information to improve the response to the next negative event. Information regarding the speed and efficiency of the response itself and how well the plan was formulated and then executed is important in improving the response. In the context of disaster recovery, the organization should learn from how well its emergency procedures protected human life and safety, and what additional procedures and equipment need to be in place for future disasters, such as fire, flood, earthquakes, tornados, and hurricanes.

Additionally, lessons learned regarding the communications process, and its effectiveness during the disaster, must be captured and should lead to improvements in communications during and after the event. There are also lessons to be learned from assessing the damage, including the qualifications of the damage assessment team, and providing for their safety during the assessment. Finally, restoring critical services, such as power, water, heat, and a facility for people to work in, is another area that must be evaluated for improvement.

All of these lessons learned, and other critical information gathered during the response, should be captured and documented as soon as possible, since relying on the memory of people may not help the organization adequately retain and use this information. As soon as the response has concluded and the organization is back to some normal or acceptable level of operations, the disaster recovery team should gather and discuss the effectiveness of the response, its successes and failures, and the lessons learned that should be used for the next disaster.

Cross-Reference

Many of the DR processes are similar to the incident management processes discussed in Objective 7.6, such as communications, training and awareness, and lessons learned.

REVIEW

Objective 7.11: Implement Disaster Recovery (DR) processes In this objective we reviewed disaster recovery processes. The major goal of disaster recovery is saving lives and preventing harm, after which the focus turns to saving equipment, systems, and data. Many of the DR processes are similar to, or even run concurrently with, incident response and business continuity processes.

Disaster response covers many key issues that must be addressed by the disaster recovery plan. The planning process should address the criteria under which a disaster will be declared and ensure the right people, processes, and resources are in place to activate the response team, assess damage, make sure the communications with all stakeholders are maintained, and then transition to business continuity activities.

The disaster response team should be comprised of experts in a variety of areas, but all should be trained in emergency procedures, damage assessment, salvage and recovery operations, and all other activities detailed in the DRP.

Communications procedures must be carefully detailed in the DRP. Communications include those that go up and down the chain of command, laterally throughout the organization, and out to external stakeholders and agencies. Organizations must identify contingencies in the event that normal communications, such as land lines, networks, and cellular capabilities, are degraded or unavailable due to the disaster. This may include using mobile radios or even physically sending people to employee residences.

Damage assessment is performed by a qualified team of knowledgeable individuals from within the organization. This team is trained and has experience in determining the survivability, repairability, and usability of equipment and systems. Damage assessments should only be performed after authorities have determined it is safe to reenter facilities and begin recovery operations.

Restoration is the process of restoring facilities, equipment, and data so that personnel are able to do their jobs safely after the disaster. Much of this effort overlaps with business continuity but must occur before the organization can get back to any functional level.

Training and awareness are crucial for organizational personnel to be prepared to execute their responsibilities during a disaster. Although classroom training is the start, exercising the DRP is how individuals will truly learn their roles and how to perform the tasks expected of them during a disaster. The importance of DRP exercises cannot be over emphasized.

Lessons learned is a concept that means the organization should take everything it has learned from the disaster recovery process and formalize it by including it in future iterations of the DRP. These lessons include those learned from resource allocation, availability of personnel, emergency procedures, recovery from equipment and facility damage, and conducting the entire recovery effort.

7.11 QUESTIONS

1. There has been a major tornado in your area that damaged the buildings of several businesses, including your company's building. Power and other utilities are out on a widespread basis, and the storm also damaged telephone lines and cellular towers. Which of the following is likely the best way to initially communicate with company personnel to determine their safety and status?
   - **A.** E-mail
   - **B.** Instant messaging
   - **C.** Runners
   - **D.** Television broadcasts
2. Public safety officials have declared your company's building safe to enter after a major fire destroyed most of the facility. Which of the following is likely the first step the organization should take toward restoration after personnel are allowed back into the facility?
   - **A.** Perform a damage assessment.
   - **B.** Relocate all personnel to an alternate processing facility.
   - **C.** Power on equipment and begin business continuity of operations.
   - **D.** Ensure personnel are trained on what they should do to assist in recovery operations.

7.11 ANSWERS

1. **C** Physically sending runners to employee residences may be the most effective way to initially communicate with them about their safety and status until other methods of communications have been restored. All of the other methods require at least power, which is out on a widespread basis, and may also require Internet access, which is likely also sporadic.
2. **A** A damage assessment is the first logical course of action to take once personnel are allowed back into the facility. This will help the organization understand what equipment, systems, and data can be recovered, and to what extent.



Objective 7.12 Test Disaster Recovery Plans (DRP)

Objective 7.11 described the various disaster recovery (DR) processes that should be included in an organization's disaster recovery plan (DRP). It also emphasized that those processes won't be effective during an actual disaster unless they are tested and practiced beforehand. This objective discusses the different techniques that your organization can use to test its DRP processes. Note that this objective also applies by extension to incident response (IR) and business continuity (BC) plans, since disaster recovery, incident response, and business continuity have closely related processes. It's not unusual for organizations to conduct response exercises that cover all three of these areas (to varying degrees), since logically, after you initially respond to an incident or natural disaster, you would recover from it (in different ways, depending upon the nature of the event), and then ensure that your business is back in operation.

Testing the Disaster Recovery Plan

It's not sufficient to simply produce disaster recovery plans and then lock them away until a disaster strikes. You really wouldn't know if those plans are effective until you had to use them, and you certainly don't want the first day of a major catastrophe to be the day that you discover that those plans don't work exactly as you thought they would. That's why testing disaster recovery plans is so critical. Every organization should test its disaster recovery plans on a periodic basis, at least annually, although more frequently is certainly desirable.

When testing the DRP, the goal is to make sure all the processes and steps work in the order they are supposed to, that all resources you plan to have for the recovery are actually available, and that alternative actions are specified in case something doesn't go according to plan. Testing and exercises are the perfect opportunities to discover and correct any issues with the plan. For example, if the plan calls for a set of specific tasks or activities to occur in sequence, the

organization might discover through an exercise that the sequence can't actually happen in real life because the resources aren't available, or that the sequence of activities simply can't happen as specified in the plan. Personnel may not be available to perform two tasks at once, for example, or someone else may be using the resources needed for a different critical task. These are the things you will discover during testing and exercises, which is why they are so important. All of these issues can be smoothed out before an actual disaster occurs.



NOTE There is a subtle difference between the terms *test* and *exercise*, even though both terms are often used interchangeably. A test usually involves determining if a particular task, such as a system cutover, actually works as planned. An exercise involves performing a programmed series of tasks that may have already been tested, to gain experience and insight in performing the overall process. Although the title of CISSP exam objective 7.12 uses the term "Test," the activities we will describe mostly involve exercising the disaster recovery and business continuity plans.

In this objective we're going to discuss the different types of tests and exercises that every organization should use for its disaster recovery plans. All of these tests and exercises are presented in order from least intrusive to normal business operations to most intrusive. Although it's easy to simply perform exercises that don't affect normal business operations, those are not true tests of what will happen during an actual disaster. Each of these exercises has its purpose, and should be used for that purpose, including testing or exercises that may detrimentally affect business operations.

Cross-Reference

Business continuity exercises using the same types of tests and exercises are discussed in additional detail in Objective 7.13.

Read-Through/Tabletop

A read-through or tabletop exercise is simply a gathering of stakeholders and participants for the purposes of going through the documented plan step by step. This helps participants become familiar with the plan and understand their general role in it. It's an opportunity for them to ask questions, and get those questions answered, about what they will be doing during recovery operations. The review can help point out obvious errors in documentation or planning, and cause people to ask questions about what they will be doing or how they will do it, what resources will be committed, and how they will get them, among other questions. Note that a read-through or tabletop doesn't have to be done with all participants physically at the table; it will likely be more effective that way, but participants can also read the documentation virtually or independently and submit questions over collaborative software asynchronously.

It's important to emphasize that a read-through or tabletop exercise is really just for familiarization purposes only. It doesn't fully exercise the plan to discover some of the practical issues associated with it. It covers the theoretical aspects of the plan; what *should* take place versus what actually *will* take place during recovery operations. Note that a read-through or tabletop exercise is generally nonintrusive to the organization's operations.



EXAM TIP A read-through or tabletop exercise is for plan familiarization purposes only; it does not actually exercise the plan. This type of test is the least intrusive of all disaster recovery tests and exercises.

Walk-Through

A walk-through exercise takes a simple plan review to the next step. In this type of exercise, participants exit the conference room, plan in hand, and walk through all the different business areas that have a role to play in disaster recovery. No actual equipment is used, and data is not transferred to alternate capabilities, but the walk-through can help people physically visualize sequences of events, places they will meet, equipment they must move, and so forth. It can help people think through the process they're going to perform, which can help them identify and point out obvious issues. For example, if servers must be pulled from a rack to be relocated to an alternate processing facility, showing people what the servers and physical space looks like might cause them to realize that they need the proper tools with them to perform the task, need to unblock access doors, need a plan to shut down servers gracefully if they are still running, and so on.

A physical walk-through is extremely helpful for people to get a better idea of how what's printed on paper will be actually implemented in the physical world. Don't be surprised if a physical walk-through test causes a great many changes and additions to the disaster recovery plan. As with a read-through or tabletop exercise, walk-through exercises normally do not affect normal business operations.

Simulation

Simulations take testing to the next logical level. As previously discussed, in read-through/tabletop exercises, participants simply review the plan, and in walk-through exercises, participants physically walk around and look at areas where they would perform tasks or activities that lead to recovery. However, during these first two types of tests, participants do not actually touch any equipment or interact with any data. Simulations allow participants to actually perform some of these tasks and activities, as well as interact with systems and data to a certain degree.

Simulations will normally be focused on specific activities, rather than an entire exercise, although the entire DR/BC process can be simulated. It's important to note that any technical activities, such as system backups or restorations, for instance, are performed on systems

that are ordinarily used as spare, testing, or backup systems. Actual systems and data in use for primary operations are not touched, since that could adversely affect the organization's actual operations.

Simulations using critical equipment for which there is no spare or backup available should be a last resort, since this could require a piece of equipment that is serving an actual operational function at that moment. To mitigate this shortfall, the organization should determine a way to simulate tasks on critical equipment using other methods, such as using mockups, virtual machines, or even software simulation programs. Simulations also should not normally interfere with the actual operations.

Parallel Testing

The previous exercises we discussed (read-through/tabletop, walk-throughs, and simulations) normally should not affect actual processing. However, the next two types of exercises, parallel testing and full interruption testing, will likely affect normal operations to some degree. A parallel test is one in which the organization actually turns on redundant or backup processing equipment or uses alternative processing capabilities and exercises them while also maintaining its fully operational capability. These two capabilities run in tandem, often using the same data or even the same systems. The purpose of this test is to determine if the alternative processing capabilities will actually function and perform their critical operations.

There is a risk that parallel testing could interfere with actual business operations, since some of the same equipment or data may be used. Additionally, exercising alternate processing capabilities at the same time the organization is using its primary capabilities may require the same personnel doing additional work and expending additional time and resources.

Full Interruption

As the name indicates, a full interruption test is the most intrusive type of exercise an organization can perform. In this type of exercise, the primary processing capabilities are completely cut over to the alternate capabilities. If there is an alternate processing site, this site is used for normal business operations for the duration of the test. This means that equipment may also have to be moved or relocated from its primary site and reconfigured to work at the alternate site. Backup or redundant systems often are used in place of primary systems to make sure that they can take on the load of critical processing. Backup data sources, such as complete copies of databases, for example, are used for processing during this type of exercise.

As much as practical, an organization should conduct a full interruption exercise as if its entire primary processing capability has been damaged or destroyed. This is the only way the organization will truly know if all of its alternate processing capabilities function as they should. In all likelihood, an exercise of this type will not be conducted for a long duration of time, since this could adversely affect real operations if something goes wrong with the cutover or if processing critical business functions using the alternate capability fails in some way.



EXAM TIP You should be well versed in the types of tests and exercises used for disaster recovery and business continuity planning. In order of least intrusive to most intrusive to business operations, these are read-through/tabletop exercises, walk-through exercises, simulations, parallel testing, and full interruption testing.

REVIEW

Objective 7.12: Test Disaster Recovery Plans (DRP) This objective addressed testing and exercising disaster recovery and other plans. Often disaster recovery and business continuity plans are exercised at the same time, using the same types of tests.

- A read-through or tabletop exercise is simply a documentation review and serves to familiarize the participants with the plan and its activities.
- A walk-through test takes the participants physically through the recovery spaces to see how they would perform their activities during a real recovery but does not involve actually performing those activities.
- A simulation test uses alternative equipment such as that used for spares, testing, or redundancy to perform tasks, helping to increase the proficiency of the participants.
- Parallel testing involves bringing up any alternate processing capabilities or alternate sites and running them side by side with primary capabilities.
- A full interruption test actually cuts over primary processing capabilities to alternate ones and tests the ability of the alternate processing capabilities to fully manage the processing load.

Note that read-through, walk-through, and simulation tests are generally nonintrusive to the organization's actual business operations Both parallel testing and full interruption tests can significantly affect actual processing operations during those tests and must be planned out carefully.

7.12 QUESTIONS

1. The business continuity planning team in your company has just completed its first draft of both the disaster recovery and business continuity plans. Since each participant has been working on their own areas and is not completely aware of the entire plan, you wish to perform a nonintrusive exercise so that everyone will become familiar with the entire plan. Which of the following would be the appropriate type of test or exercise for this effort?
   - **A.** Simulation
   - **B.** Parallel test
   - **C.** Full interruption test
   - **D.** Read-through/tabletop exercise

2. Your company has been working on its disaster recovery and business continuity plans for some time and has finalized all of its processes and activities, as well as developed its alternate processing capabilities. However, no one is sure that the alternate capabilities, when actually turned on, will function as they are supposed to. The organization wants to test these capabilities, but at the same time does not want to run the risk of shutting down all actual operations. Which of the following is the most appropriate test to perform to meet these requirements?
   - **A.** Walk-through exercise
   - **B.** Parallel test
   - **C.** Simulation
   - **D.** Read-through/tabletop exercise

7.12 ANSWERS

1. **D** Since the goal is to have participants become familiar with the entire plan but without intruding on business operations, a read-through or tabletop exercise is the best choice. It is nonintrusive, does not require participants to perform any procedures with which they are not yet familiar, and does not use actual equipment.
2. **B** A parallel test allows the organization to exercise its alternate processing capabilities without having to shut down its primary operations. Although this is somewhat intrusive, the alternate processing capabilities must be tested at some point before an actual disaster strikes.



Objective 7.13 Participate in Business Continuity (BC) planning and exercises

As a reminder, the basics of business continuity planning were discussed in Objective 1.8. Specifically, we discussed the first key step of business continuity planning, the business impact analysis. Objective 7.13 closes our discussion on business continuity. In this objective we will discuss the process involved in completing the business continuity plan and then review various business continuity exercises, which are similar to the disaster recovery plan exercises introduced in Objective 7.12.

Business Continuity

As noted in previous objectives, business continuity planning is a separate endeavor from disaster recovery planning, although they are closely related. Think of how negative events happen in a sequence and what takes place during that sequence. An overall incident response is

what the organization must do first to triage and contain the incident, whether it is a malicious attack or a physical event, such as a fire, tornado, and so on. Although we tend to think of incident response as applying only to information systems that are the target of a malicious human attack, that is not always correct. Then comes the disaster recovery effort (if the incident threatens lives, safety, equipment, or facilities), and then the last piece is business continuity—getting the business back in operation.

However, incident response planning, as well as disaster recovery planning and business continuity planning often occur in parallel. Business continuity planning focuses on ensuring that the business can still function at some operational level, even if it is not at the optimum level. Disaster recovery planning, on the other hand, is concerned with ensuring safety and protecting human life first, with a secondary goal of preserving equipment and facilities.

Although much time and effort are invested in disaster recovery planning, when a disaster actually strikes, implementing the DRP is a reactionary process based on the type of disaster and other factors. Once the incident response team determines that a disaster has occurred, disaster recovery takes over so that human safety is preserved and, subsequently, attempts are made to preserve equipment and facilities. Once those things are preserved or salvaged, business continuity begins so the organization can resume operations within an acceptable timeframe and at an acceptable level.

Business Continuity Planning

The goal of business continuity planning is to ensure that the organization can restore some level of operations as soon as possible after an incident or disaster. Ideally, the organization would reach a fully operational status quickly, but this isn't always possible. Several factors impact the pace at which an organization can achieve an operational level, including whether and when necessary resources are available after a disaster and the level of damage to the infrastructure. To achieve an acceptable level of operations after an incident or disaster, the organization must understand the intricacies of all of its business processes, particularly those that are critical to its operations.

One of the first things the organization must do to begin the business continuity planning process, assuming it is starting from scratch in this process, is to staff a business continuity planning team with qualified members of the organization. The business continuity planning team will primarily be business process owners, senior managers, and a variety of other advisors from functional areas. IT and cybersecurity must be part of the team but will not be the focus of the team, except to the extent that they are working to preserve crucial information resources that support those critical business processes. In addition to IT and cybersecurity representatives, representatives from other functional areas may include accounting, facilities and personnel security, human resources, and other support areas.

Business Impact Analysis

Although we already discussed the BIA process earlier in the book in Objective 1.8, it bears repeating here in this discussion on business continuity planning. Recall that the BIA attempts to inventory and prioritize the critical business processes that make up the organization's

operations. This also necessarily means inventorying all the critical assets, such as systems and information, that support those critical processes. These critical assets are the primary concern of the business continuity planning team.

A disaster may damage or destroy many systems, pieces of equipment, facilities, and information. The question then becomes how to recover, repair, or replace those assets or make adjustments so those critical business processes are still supported. The BIA process focuses on the priority each of these assets has for restoration; the result of this process is a detailed analysis of these assets and how long the organization can function without them, which then prescribes their priority for restoration, and how the organization might go about doing exactly that. The BIA gives the organization clear direction on which critical processes and assets they must focus on for the business continuity plan itself.

Cross-Reference

Business impact analysis was discussed in more depth in Objective 1.8.

Developing the Business Continuity Plan

The business continuity plan details the processes and activities that must take place during and immediately following a catastrophic event, so that the organization can begin to work again. There will likely be minimum requirements that the business must have in place after a disaster to even begin functioning; these usually include a stable processing facility with power, water, heat, communications, and other necessities. They also include the proper functional equipment, such as computing systems, and the most current information available to begin processing. This information might be in the form of customer order databases, accounting software, basic office productivity suites, and so on. The organization must plan on ensuring the availability of these resources down to the lowest level of detail in order to account for everything the organization will need to begin functioning again.

Important Business Continuity Metrics

An organization's business continuity plan should, to the maximum extent possible, be very precise in laying out the details of how long the organization can truly function, assuming some of its capability has been damaged or destroyed. These details help determine the prioritization for restoration of critical processes and the resources that support them, as well as prioritization for actions that need to take place during the recovery process.

Some key business continuity metrics that you should know for the exam include

- **Maximum tolerable downtime (MTD)** This is the maximum amount of time that a business process or function can endure disruption before suffering catastrophic harm.
- **Recovery time objective (RTO)** This is the maximum amount of time the organization can tolerate a disruption to a process or system.

- **Recovery point objective (RPO)** This is the maximum amount of data, measured in time (e.g., two days' worth of data), that the organization can afford to lose; it is also the minimum amount of data that the organization must attempt to recover if needed.
- **Mean time to failure (MTTF)** This is the amount of time a component or system is rated to function before it eventually fails. This is generally the entire lifespan of the component if it cannot be repaired.
- **Mean time between failures (MTBF)** This metric applies to components that can be repaired; this is the estimated amount of time between failures. The organization should use this metric to determine how many spares it should keep on hand for system components.
- **Mean time to repair (MTTR)** This indicates the time needed to repair a damaged component; it may include not only the actual time to complete the repairs, but also the amount of time needed to obtain parts needed for the repair through the supply chain.



EXAM TIP You should be familiar with key business continuity metrics, particularly recovery point objective and recovery time objective. Remember that recovery point objective is the amount of data, measured in time, that the organization can afford to lose before recovery becomes difficult. The recovery time objective is the amount of time the organization can tolerate to recover a particular business function.

Business Continuity Exercises

We discussed the different types of exercises in Objective 7.12; as a quick review, these include

- **Read-through/tabletop** Familiarizes participants with the steps of the plan and identifies obvious issues or shortfalls in the plan; nonintrusive to actual operations
- **Walk-through** Physically steps participants through the plan without actually performing any activities or using any resources, such as systems or information; also nonintrusive
- **Simulation** Uses nonproduction equipment and data to simulate the tasks that must be performed during the actual event; does not interfere in actual operations
- **Parallel testing** Uses redundant equipment and capabilities to test actual activities while simultaneously keeping normal operational capabilities running; can be intrusive on actual operations due to the commitment of people and resources
- **Full interruption testing** Involves a full cutover from primary operational capabilities to alternative operations; completely interrupts normal processing to test the disaster recovery and business continuity processes

Cross-Reference

The types of tests and exercises that are used for both disaster recovery and business continuity were discussed in Objective 7.12.

Understand that since disaster recovery and business continuity are so closely linked, these types of exercises are also used to exercise the BC plan. In fact, it is not unusual to see the full gamut of incident response, disaster recovery, and business continuity plans combined for each of these types of exercises. The key point you should remember about these exercises is that they test the viability of the disaster recovery and business continuity plans. This means that, with each exercise, both of these plans should improve increasingly. Shortfalls in the plans are usually discovered during exercises, and they are changed, tweaked, or even completely redone in some cases. It's necessary to exercise the plans on a frequent, periodic basis. This is because the operating environment changes and people leave and join the organization, so levels of expertise and experience change. Exercising the plans ensures that new employees become versed in the disaster recovery and business continuity processes, and people who are already part of the organization get better at executing the plans.

Scenarios used during business continuity exercises should be realistic and test all aspects of the business continuity plan. This includes transferring functions to alternate processing capabilities or sites, simulation of downed equipment that must be replaced or repaired, reconfiguring services to support critical business processes where needed, and coming up with alternate methods to process information in case assets cannot be repaired or replaced.

REVIEW

Objective 7.13: Participate in Business Continuity (BC) planning and exercises This objective completes our discussion of the overall disaster recovery and business continuity processes. In this objective we discussed the business continuity planning process, including the importance of the planning team itself. We briefly described a few of the important metrics that help determine how long the business can function, including metrics related to component failure and repair, as well as the recovery time objective and recovery point objectives. We also reiterated the importance of the business impact analysis process and the importance of participating in business continuity exercises. These exercises are often combined with disaster recovery exercises.

7.13 QUESTIONS

1. Which of the following is the first and likely most important part of the business continuity planning process?
   - **A.** Business continuity exercises
   - **B.** Business impact analysis

C. Developing business continuity metrics

D. Assembling the business continuity planning team

2. Which of the following business continuity metrics represents the maximum amount of data the organization can afford to lose, in terms of time, before its successful recovery is threatened?

A. Maximum tolerable downtime (MTD)

B. Recovery time objective (RTO)

C. Recovery point objective (RPO)

D. Mean time to failure (MTTF)

7.13 ANSWERS

1. **B** The business impact analysis is the first part of the business continuity process, and likely the most important aspect since it is during this part of the planning that critical business processes and assets that support them are determined.
2. **C** The recovery point objective represents the maximum amount of data that an organization can afford to lose before its full recovery becomes unattainable; it is measured in terms of time.



Objective 7.14 Implement and manage physical security

We have discussed the importance of physical security controls in general in other domains and objectives, but in this objective we will focus on specific physical security controls, both perimeter security controls that protect facilities, equipment, and people, and internal security controls designed to mitigate insider threats and unauthorized access.

Physical Security

Cybersecurity personnel often assume that the responsibility for physical security belongs to another group of people, skill set, or security program. However, as a candidate for CISSP certification, you must understand the importance of physical security and how physical security controls integrate with administrative and logical security controls. Firewalls, intrusion detection systems, rule sets, strong authentication, encryption mechanisms, and other administrative and logical security controls will not protect your organization's vital information if an intruder has physical access to a system or piece of equipment. An old adage in security

states that if a malicious individual has physical access to your computer, they now own your computer. Physical security controls are often the first line of defense against harm to individuals, equipment, and facilities, and the importance of such controls to cybersecurity and CISSP candidates cannot be overstated.

In this objective we will discuss physical perimeter security controls such as security guards, controlled entry points, fencing, proper lighting, and barriers. We will also discuss physical internal security controls such as security zones and physical security processes and procedures. This discussion completes and supplements what we have already discussed in Domains 3 and 5.

Cross-Reference

Physical security was also discussed extensively in Objectives 3.8, 3.9, and 5.1.

Perimeter Security Controls

An organization's logical perimeter can be more difficult to define than its physical perimeter, since there are so many logical entry points into the organization's network infrastructure—through firewalls, over wireless connections, and even through applications that touch the Internet, like e-mail. Although it may be easier to define, the physical perimeter of an organization often is just as difficult to protect. The complexity of physical perimeter security varies depending on several factors, including how the facility is laid out, the number of people that work in the facility, the ebb and flow of foot traffic entering and exiting the facility during particular times of day, the number of security guards available, the types of gates and fencing, and even the lighting around the facility. In this section we are going to discuss several aspects of perimeter security controls, including video surveillance, barriers, fencing, security guards, and physical intrusion detection systems.

Recall from Objective 3.9 the discussion of the discipline of Crime Prevention Through Environmental Design (CPTED), which promulgates four primary areas that should be considered in designing perimeter controls: natural access control, natural surveillance, territorial reinforcement, and site and facility maintenance. You should evaluate these considerations when designing and implementing perimeter controls such as barriers, fencing, traffic flows, and other aspects of perimeter security.

Perimeter Security Zones

Physical security zones serve the same purpose as logical security zones. Remember from Objective 4.2 that logical security zones are created between untrusted networks and sensitive networks to establish a barrier between external and internal networks (e.g., demilitarized zones created by firewalls and border routers). Physical security zones are also created to separate untrusted or public areas from sensitive areas, such as those that house sensitive equipment and information processing activities. Security zones can be employed around the perimeter of the facility to separate publicly accessible areas from areas that only employees

can access, such as employee parking lots, areas inside a security fence, outdoor break areas, and so on. Security zones can be delineated by fences; obstacles, such as bollards; or natural objects, such as shrubbery. Security zones outside of a facility should be monitored by security cameras, guards, and by organizational personnel. Security zones can also be used inside the facility, as we will discuss later in the objective.

Entry Control Points

Entry control points are critical in controlling the flow of personnel into a facility. Small to medium facilities should have a single primary entry control point where employees, guests, business partners, contractors, and other legitimate personnel are allowed to enter and then have their identification checked and validated. Larger facilities could have more than one primary entry control point. There may be secondary entry control points for employee-only entrances, as well as entrances opened temporarily for deliveries, shipments, large equipment, and so on. These secondary entry control points should also be carefully controlled and monitored.

Based on the security requirements of the facility, entry control points may be automated, but a security guard should be present to handle unexpected issues and video surveillance capability. Unauthorized entry by vehicles or pedestrian traffic can be prevented through the use of barriers when required and practical. If possible, entry control points should be located a reasonable distance away from the main facility, such that the distance will help delay or deter someone who gets past the entry control point by crashing a vehicle through it or running through the barriers. The objective is to give the guard enough time to alert security patrols or law enforcement, if possible, before the intruder gets to the main facility.

Fencing

Fences establish the physical boundaries around an organization's facility. They are used to distinguish which parts of a property are controlled by the organization and which are not. There is usually an exterior perimeter fence that surrounds the campus of an organization, but there also may be interior areas that are fenced in, if those areas require higher security. Use of multiple fences helps add to layers of physical security.

Fencing comes in various sizes and construction materials:

- Fences under four feet are usually sufficient to keep casual intruders out and establish a boundary around a sensitive area of a facility.
- Fences up to eight feet high will deter most trespassers.
- Fences over eight feet high are designed for higher security areas.

Sometimes fences have barbed or razor wire (called concertina wire) inserted into spirals at the top of the fencing to add a layer of difficulty for anyone who is determined to climb the fence.

A Perimeter Intrusion Detection and Assessment System (PIDAS) is a type of fencing system that may have two or more deterrents used together to create physical security zones.

You'll normally see a PIDAS used for secure military installations and critical infrastructure, such as power plants and electricity substations. The inner fence may be over eight feet tall and be electrified or have concertina wire across the top. An outer fence may be much shorter and less secure, intended to keep out people who are simply curious, or even wild animals. Depending on the security level required by the installation, additional layers of fencing may be needed. Between fencing layers, there is typically an area for security guards to patrol, which can also serve to hold any intruders who penetrate the outer fence in a segregated area. This type of fencing is usually also accompanied by intense flood lighting mounted at the top of tall poles that can illuminate the entire area. PIDAS implementations also usually have physical intrusion detection sensors and alarms to notify security personnel if the inner and outer fences are breached.

Fencing works hand-in-hand with gates and turnstiles. Gates are used to create controlled entry and exit points. They can be as simple as a hinged piece of fencing that opens and closes or a guard house where personnel must enter for authorization before being allowed to proceed inside the fence. Many entry control points hold turnstiles, which are doors or barriers that only rotate or open in one direction, to allow personnel to either enter or exit, but not both simultaneously.

Barriers

Barriers are used to deter and delay intruders and can consist of a wide variety of construction materials, including steel, concrete, and other heavy materials. Barriers are typically deployed inside the perimeter area of a facility to route and control vehicle traffic and to prevent errant vehicles from getting too close to a building. Barriers can also be more subtle, in the form of planters, posts, and carefully planned landscaping to provide a natural-looking but carefully controlled path for pedestrian traffic. Examples of barriers placed around the facility include concrete blocks and bollards, which are effective at rerouting or blocking both pedestrian traffic as well as vehicle traffic.

Lighting

Exterior lighting is useful as a deterrent control, as trespassers, intruders, or even curious passersby are less likely to enter a facility at night if it is brightly lit. Exterior lighting also serves to illuminate the immediate area to prevent crime, such as personal attacks, theft, and destruction of property.



EXAM TIP Make sure you should understand the term *lumen*, which is the measure of light brightness or intensity. A lumen, which is described as a foot-candle (intensity of light over distance), is used to set standards for exterior lighting. A lumen illuminates an area of approximately one square meter, or approximately three square feet (called a *lux* in that instance). Security areas should normally be illuminated at an intensity of at least two lumens (abbreviated as *lm*).

Placement of light poles is also a consideration in exterior lighting. Poles should be placed close enough together to prevent dark areas between poles. Essentially, if the light from a pole covers an area of 50 feet, then the poles should be placed no more than 50 feet apart. Overlapping coverage of light from each pole minimizes unlit areas between the poles.

Surveillance

Most surveillance today within a secure facility is accomplished using video or closed-circuit television (CCTV) cameras. Video surveillance serves two important control functions:

- **Detection** It records events visually, so it can be used as a detective control.
- **Deterrence** If the video camera is visible, it can deter people from performing illegal acts or policy violations.

Surveillance is also performed by using human guards, who can react much faster to situations that they see, but are less effective than video surveillance in terms of detailing exactly what happened. Plus, being able to play back the video is useful for investigations.

Modern surveillance systems are built into the network system and use IP-based cameras, rather than the older CCTV cameras. Modern cameras also have the ability to record for long periods of time and store the data on removable media, such as an SD card. Sophisticated cameras can be controlled by a guard at a remote location; modern cameras can circle around and view an entire room or, in the parlance of video, pan, tilt, and zoom (PTZ). Advanced cameras can detect objects in motion and can perform facial recognition. Cameras can be a key component of a larger intrusion detection system that can trigger electronic alarms, initiate automated security procedures, and notify security personnel.

Guards and Dogs

Even the best automated systems, including state-of-the-art intrusion detection systems, alarms, video surveillance, and badging systems, cannot completely replace human security guards. These technologies should supplement human security guards, not replace them. Technology can help be the eyes and ears of guards, as they can't be everywhere, but an organization should always have humans present to control security mechanisms and processes. Humans can make rapid decisions that automation cannot and can decide whether or not an event is a false alarm or an issue that needs to be escalated. They can determine whether an emergency or crisis situation is valid or not, and they can make split-second decisions on what to do next to protect lives, facilities, and equipment.

There are advantages and disadvantages to using human security guards. The advantages are exactly as just discussed—their ability to make proper decisions based on context and to inject experience, training, and common sense into a situation to determine whether it is a valid emergency situation. Humans can override an automation system if the system fails to detect an event, generates a false positive, or is making an improper access control decision. Human guards also serve to build a bridge between employees and security functions, and

they are a valuable resource to help ordinary employees navigate complex security mechanisms to gain authorized access to facilities.

The disadvantages of human guards are that they can be expensive. The cost in terms of salaries, benefits, and supervision can add complexity to the human resource management process. Guards must be cleared for the areas they patrol, and they must be trained to adequately make critical split-second decisions when needed. They also add a (necessary) layer of complexity to security controls. As with all employees, guards have the typical personnel issues, such as personality conflicts, discipline, development, and so on. For smaller organizations that cannot afford to employ guards, receptionists or other employees may have to be trained on physical security measures and be utilized in areas outside their areas of expertise, such as visitor control, physical surveillance, and so on. Large organizations, however, will find a human guard force is almost an absolute necessity.

Guard dogs can be problematic. While they do have their uses, typically to patrol the outside of a facility within a perimeter fence area, they can also present issues. These include adequate training and control of the dogs, the costs of veterinary care, housing, and feeding, and, of course, liability. If a dog harms an employee who is carrying out an authorized duty, for example, the organization may be liable. An organization must carefully weigh the benefit of using guard dogs for only specific functions and in certain areas versus the cost of care and supervision and the potential liability if the organization were to be sued for harming an individual.

Internal Security Controls

Internal security controls are no less important than perimeter security controls. The goal of perimeter security controls is to ensure that only authorized individuals enter a facility; however, individuals who are authorized to enter a facility, such as guests, vendors, contractors, and business partners, are not necessarily trustworthy. Once an individual is inside the facility, even if they have been authorized to enter, they may not have authorization to access certain sensitive work areas where sensitive systems and data are located.

Once a person has been vetted and allowed past the facility perimeter (considered a visitor or authorized guest at this point), internal security controls are meant to keep them away from sensitive areas for which they are not authorized. Even in cases where there are authorized personnel accessing sensitive areas, internal security controls can help prevent unauthorized access and enforce the security principles of least privilege, separation of duties, and need-to-know. Internal security controls also help protect individuals, equipment, and facilities from malicious entities who wish to do them harm. This includes both insider threats and intruders that may specifically enter a facility to perpetrate theft, fraud, or physical harm.

Cross-Reference

Designing facility security controls was discussed in Objective 3.9 and covers internal facility areas of concern, including wiring closets, server rooms and data centers, and sensitive work areas.

Earlier in this objective we discussed the use of security zones around the perimeter of the facility. An example of an internal physical security zone would be the reception area of a facility, where the public may enter to gain initial access to a business, show their identification, and then be escorted to other areas. Past the reception area would be common areas where only employees are allowed, such as offices, break rooms, restrooms, and so on. These employees would have to meet the basic requirements of background checks and have access to these areas based on their job requirements. Other areas deemed more sensitive would only be accessible to employees who have the proper clearance, need to know, and position requirements. These sensitive areas would not be accessible to all employees. So, essentially, this example has three internal security zones in the facility. This is a simplistic view, of course, and security zones can be separated out in more detail as well.

Personnel Entry Requirements

Everyone who enters a facility, especially employees, must be there for a valid reason and be appropriately cleared. Visitors, potential employees, contractors, customers, business partners, and so on aren't considered employees but, obviously, should also have some valid reason for simply entering the least secure zone of the facility, which may be the reception area, vehicle loading zone, checkpoint, or even the initial entry control point. These individuals should be escorted at all times or should have a restricted or temporary badge that grants them limited access to an area (such as the break room in the case of a vending machine contractor), often unescorted but still controlled.

Personnel such as employees of permanent contractors and business partners should have the proper background checks, need-to-know clearance based on job function, and supervisory approval before being granted access to facilities and the sensitive areas within. Access should be only granted to specific areas, and then should be periodically reviewed to ensure the individual still has the same access needs.

Locks

It's often said that locks only keep out honest people; this may have been true with older manual locks, which could be effectively bypassed using a crowbar or a few simple tools, such as a screwdriver or lock picking kit, but modern locks, particularly electronic ones, are not so easy to circumvent. Locks are useful to deter and delay intruders and come in several varieties, including both mechanical and electronic locks. Mechanical locks range from padlocks and deadbolts to mechanical key locks and combination locks. These locks can be useful as delaying mechanisms, but their ability to delay depends upon how long it takes an intruder to compromise the lock by breaking it or picking it.

Mechanical locks come in three varieties:

- **Warded lock** A good example of a warded lock is a padlock, which consists of a spring-loaded bolt that has a notch cut in it. The key is inserted into the notch and slides the bolt when turned from the locked position to the unlocked position. The wards are metal protrusions that must match corresponding notches on the key in order for the key to be inserted and turned.

- **Tumbler lock** A tumbler lock is slightly more complex than a warded lock. The metal key fits into a cylinder and is turned, causing metal components inside the lock to move into the correct position so the bolt can be moved to either the locked position or the unlocked position. The three types of tumbler lock include the pin tumbler, wafer tumbler, and lever tumbler.
- **Combination lock** A combination lock has internal wheels that must line up correctly for the lock to disengage and become unlocked.

Electronic locks, also known as *cipher locks*, can be programmed. This type of lock requires a specific combination to be entered into a keypad or requires a badge with an electronic strip to be swiped through a badge reader or placed near its sensor (called a *proximity badge*). Cipher lock combinations can be changed, and cipher locks can also have multiple combinations so that different users can be assigned unique combinations and be audited for their access. Cipher locks also have the following security features:

- **Door delay** This feature triggers a physical alarm or alerts security guards if a door is held open too long.
- **Key override** A specific combination can be programmed into the lock to override normal functions in the event of an emergency situation.
- **Master key** This function enables security personnel to change a user's access code and program the lock.
- **Duress code** This allows an individual who may be under duress or coercion to input a secret code that alerts security personnel.

Note that electronic cipher locks also have other advanced features, such as the ability to assign specific combinations to users for auditing and access control, as well as granular control of an authorized user's ability to access the facility, such as programming the lock to work only for specific times of the day.

Physical Intrusion Detection Systems

Physical intrusion detection systems (IDSs) are used to detect illegal access or violations of security policy, as well as alert security personnel to those events. Several types of IDSs are commonly used, and they have a variety of sensors to detect anomalous events. IDSs not only detect intrusions through doors, windows, fences, and other unauthorized entry points, but can also detect theft of equipment. Physical IDSs detect

- Changes in light, through *photoelectric* or *photometric* sensors
- Sounds and vibrations, through *acoustical detection systems*
- Changes in temperature, through *passive infrared systems*

- Motion, by using *wave pattern motion detectors*, which sense disturbances in the patterns of microwaves, ultrasonic waves, or low-frequency energy waves, and *proximity* or *capacitance detectors*, which measure disruptions in an electromagnetic field and can be set up around an object, like a piece of artwork, to detect theft
- Changes in electric circuit continuity through *electromechanical systems* (useful for detecting doors or windows that have been opened)

All these changes to environmental factors can indicate the presence of an intruder, provided the right sensors are in place.

The following are several other things you should understand about physical IDSs:

- They can be expensive, depending on what type of IDS you get and how sensitive it is to the environment.
- They should have redundant power supplies or emergency backup power in case of a power failure.
- They should be linked to a centralized, integrated physical security system. The fail-safe configuration of these systems should default to activated, so that in the event of a power failure or some other interruption of service, the IDS remains turned on.
- They should alert security personnel if any attempt is made to tamper with them.
- They require human monitoring and interaction when they trigger an alarm.

Auditing Facility Access

Like access to information systems, access to facilities, particularly sensitive areas, must be audited. You should audit facility access frequently and routinely to detect trends or specific instances of abnormal behaviors. This includes auditing not only strange events, but also the massive number of allowed routine events to detect possible negative trends. Note that you will have to filter through a lot of "noise" (innocuous events) and be able to correlate seemingly unrelated events over time to determine if there are any trends or particular events of interest.

Auditing facility access typically involves a combination of manual and automated processes. Manual methods of recording access control, such as a pen-and-paper visitor's log, should be reviewed daily and transferred to a database system whenever possible. At a minimum, manual logs should be kept a specific period of time so that they can be reviewed later in the event they are needed for an investigation. They are also correlated with automated logging methods to support timelines, events, and security processes.

Automated means of recording facility access mirror those that are used for system access; normally a person's identification information (such as badge number), type of event (entry to or exit from an area), and date and time of the event are all recorded through badge sensors, video surveillance and facial recognition, thumbprint scanners, and so on.

Often data from automated methods is collected and shipped to a centralized logging facility, which may also include system access data and other sensor data. Combining logs can create a complete picture of an individual's activities from the time they entered a facility—where they went, what systems they accessed, and the interactions they had with systems and data. All this data is correlated and analyzed through a security information and event management (SIEM) system and can be used to accurately pinpoint everything an individual does inside the facility.

You should audit for specific events that trigger alarms and notify security personnel of failed entry or exit attempts and tailgating (also known as *piggybacking*, where someone closely follows an authorized person through an entry or exit point without being authenticated). Also closely monitor access granted temporarily to authorized guests, vendors, partners, and so on. This applies not only to the common areas of the facility but also to all sensitive work and processing areas.

REVIEW

Objective 7.14: Implement and manage physical security In this objective we continued our discussion of physical security by drilling down into categories of perimeter security controls and internal security controls.

Perimeter security design should consider the following security controls:

- Security zones are used to separate sensitive areas both in the perimeter of the facility, as well as internal sensitive areas
- Entry control points into a facility should be single points of entry or exit where personnel traffic can be controlled.
- Fencing is used to protect the perimeter of a facility and can be various heights and have features such as concertina wiring. Multiple fences may be used to create physical security zones.
- Barriers are used to block and route both vehicle traffic and pedestrian traffic and can be natural barriers or made of concrete, steel, and other materials placed in key areas around the facility.
- Both exterior and interior lighting should be bright enough to illuminate the area and prevent someone from being able to hide or remain undetected.
- External lighting should be spaced so that there are no dark spots between light poles.
- Surveillance can be human-based, in the form of security guards, or technology-based, in the form of video cameras.
- Human guards provide an added benefit in that they can react quickly and make split-second decisions when needed to protect lives, equipment, and facilities.
- Guard dogs can be used in limited situations, but often incur expense and liability.

In addition to controls mentioned for the perimeter, internal security controls that should be considered include the following:

- Personnel entry requirements usually are guided by security clearance, need-to-know based on job position, supervisory approval, and reason for access.
- Locks can be manual or electronic and provide an additional layer of protection for sensitive areas.
- Auditing facility access is necessary to ensure that only authorized personnel enter the facility and sensitive areas. Both success and failure of physical access events should be audited to determine both trends and to see significant events.

7.14 QUESTIONS

1. You have been asked to be part of a security committee that is tasked with developing recommendations for improvements in facility security. One issue that has been brought up is the fact that people can enter the facility through any of four entrances. Which of the following suggestions would mitigate this issue?
   A. Eight-foot fencing with concertina wire
   B. Security zones
   C. Centralized single entry control point
   D. Guard dogs
2. You are designing security for the perimeter of a secure data center and need to determine floodlight spacing around the exterior perimeter fence, which is about 140 feet long on each side of the facility. The type of floodlights you've selected cover an area of about 70 feet in diameter. Based on this, how far apart should the light poles be installed?
   A. 50 feet
   B. 70 feet
   C. 140 feet
   D. 100 feet

7.14 ANSWERS

1. **C** Installing a centralized single entry and exit point for the facility would prevent people from coming from all directions off the street into the main entrance. Although fencing would help achieve this goal, by itself fencing would not consolidate entry and exit of personnel into a single point. Guard dogs would not help limit entry into the facility and would introduce too much liability for the organization.
2. **B** The light poles should be installed so that there are no unlit areas between them. If the effective diameter of lighted area from a pole is about 70 feet, then the poles should be about 70 feet apart, so that each lighted area borders the next area for the next pole.



Objective 7.15 Address personnel safety and security concerns

In this objective we will discuss personnel safety and security. Many aspects of this topic have been covered in other objectives as well, but they bear repeating here because the preservation of human life and the safety and security of personnel are the most important priorities within the organization. We will discuss personnel safety and security within the contexts of travel, security awareness and training, emergency management, and duress.

Personnel Safety and Security

Giving precedence to the preservation of human life and the safety and security of personnel is required not only by numerous laws and regulations but also to imbue the organization with high moral and ethical standards. People are the most important resource the organization has; buildings, equipment, and even information systems and data can be replaced, but human lives cannot. We discussed personnel security in other objectives, so we will focus primarily on the safety aspect of this objective.

Travel

Before an employee travels on behalf of the organization, especially an executive or an employee carrying sensitive information of the organization, they need to be made aware of the broad range of threat actors who specifically target business travelers. They may not necessarily be a target due to where they work or what information they have, but anyone can fall victim to criminal activities. In addition to the precautions that any reasonable person should take in a potentially harmful situation, employees who have sensitive information or have critical responsibilities should be careful when traveling to prevent harm or compromise to themselves and to company assets, such as mobile devices.

Some of the specific precautions employees should take when traveling include

- Take a mobile device containing your organization's sensitive information only if absolutely necessary; in that scenario, minimize the sensitive information to only that which is needed on the trip and ensure that it is encrypted.
- Take special precautions to maintain physical control of said mobile devices. During the times when maintaining physical control is impossible, such as when traveling through airport security, ensure the device is locked or powered down.
- Be aware of your surroundings at all times.

- Have a plan when traveling, particularly to high-risk international areas or areas with higher crime rates and share it with people who are traveling with you or at least people back at your home. Examples include specific times that you will call them or meet, as well as alternate meeting sites and what they should do if they don't hear from you within a certain amount of time.
- Avoid free Wi-Fi hotspots, unless they are in a trusted location.
- Only use secure virtual private network (VPN) connections to access company network resources, including e-mail.

Security Training and Awareness

Although discussed previously, it's important to emphasize the concept of security training and awareness, particularly when discussing personnel security and safety. We already discussed the differences between awareness, training, and education in Objective 1.13. Rather than review the types of training available or presentation techniques you should provide for your personnel, we will focus on the content of the training, and in this context, it should concern their own personal safety and security.

You should design and implement formal training programs that help employees increase their safety awareness and develop improved security habits. Cyber hygiene is a good topic to train on in terms of general personal security. It can help personnel not only understand the criticality of obeying security policies at work but also develop cyber hygiene habits that will protect their own personal information at home.

Security awareness and training focused on personnel security and safety generally includes the following topics:

- Personal and professional cyber hygiene (e.g., control of personal information, awareness of social engineering techniques, control of portable storage devices, etc.)
- Awareness of surroundings and environment
- Common sense precautions to take when traveling
- Adherence to emergency procedures, such as fire evacuation plans and active shooter protocols
- General workplace safety procedures

Cross-Reference

Security training and awareness were also discussed in Objective 1.13.

Emergency Management

Emergency management is a part of the organization's processes for keeping its personnel safe from harm and should detail their due diligence and due care responsibilities. Emergency management includes obvious processes, such as having a fire evacuation plan, workplace

safety regulations, and accident reporting requirements, but it involves much more. You could consider disaster recovery, for example, as part of emergency management, since disaster recovery planning is also concerned with saving lives and reacting to disasters.

The organization must, at a minimum, comply with health and safety regulations. Those include having emergency evacuation plans, placing the proper safety equipment throughout the facility, and ensuring that its personnel are trained to use that equipment. The organization should also have an emergency management point of contact, preferably someone in senior management, who is responsible and accountable for ensuring that the emergency management team implements and maintains these processes.

At a minimum, the organization should ensure that it has in place the following emergency management processes and procedures:

- Emergency evacuation plans
- A formalized safety program
- Recurring safety training
- Safety-related risk assessments
- Safety and emergency equipment prepositioned at various critical points within the facility
- Fire prevention, detection, and suppression systems
- Environmental monitoring systems
- Disaster recovery procedures

Obviously, the organization should test these processes and procedures on a regular basis so that personnel, especially newly hired employees, know what to do in the event of an emergency or other crisis.



EXAM TIP Keep in mind the various steps the organization should take in implementing its emergency management processes, to include evacuation plans, safety training, installation of safety equipment, and so on.

Duress

A person under *duress* is either being physically harmed or being threatened with violence or some other harmful action in order to make them do something against their will. Since people are the most valuable, and typically vulnerable, asset in an organization, they can often be threatened with violence to coerce them to do something that may harm them, other people, or an organization's information systems. An example might be if someone were coerced with the threat of violence against their person to give up an administrator password. Although this may sound like the stuff of movies, these events actually do happen. Most of the time, individuals are alone or in a vulnerable situation when they are placed under duress. Think of an

administrator who is working at night and may come out to a deserted, dark parking lot where a disgruntled employee may accost them, for instance.

Duress systems and procedures are used to alert others who can help. Consider a bank teller who is able to press a secret button to notify law enforcement in the event of a robbery. Duress systems can be placed in a data center, reception area, or sensitive processing area, in case an intruder is able to enter the facility. Duress systems could alert guards or law enforcement, automatically lock doors, and sound alarms. Duress procedures can also assist if a person is unable to engage a duress system. For instance, if an employee is being threatened while on the phone or talking to a coworker to obtain information or to give orders to that person, they could use a specific code word or phrase that is known to both of them as the organization's duress signal. When the coworker hears this word or phrase spoken by the employee in danger, they could then take action to alert the authorities.

REVIEW

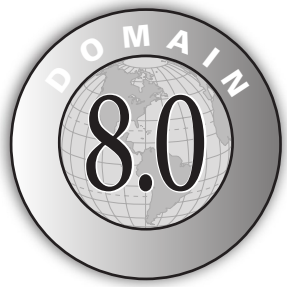
Objective 7.15: Address personnel safety and security concerns In this last objective for Domain 7, we discussed the importance of personnel safety and security, which should be the highest priority in an organization. We touched on aspects of several topics, including security training, safety awareness, and emergency management. We also discussed the importance of employee safety during travel as well as notification options if they are under duress.

7.15 QUESTIONS

1. Your company's security awareness and training program is comprehensive in its coverage of cyberthreats. Which of the following topics should you also include in the program to address issues of personnel safety and security?
   - **A.** Social engineering
   - **B.** Physical intrusion
   - **C.** Malware
   - **D.** Phishing attacks
2. Your company has hired an outside consultant to train its personnel on emergency procedures that might involve a hostage situation. The scenario that the consultant has given your employees involves not being able to trigger a physical alarm to alert security personnel. Which of the following is a useful technique to alert someone without having to trigger a physical alarm?
   - **A.** Using a duress code word
   - **B.** Yelling to security guards that may be within earshot
   - **C.** Attempting to place a call to the security office when the opportunity is afforded during the situation
   - **D.** Attempting to escape to alert someone about the situation

7.15 ANSWERS

1. **B** A physical intrusion could directly affect personnel safety and security, particularly if the intruder begins an attack of a violent nature, such as that instigated by an active shooter situation. The other options are not as relevant to personnel safety or security.
2. **A** Using a duress code word should be considered when in a situation where triggering a physical alarm is not possible or safe. The other options mentioned could result in physical harm to individuals and should not be considered unless there is no other option.



Software Development Security

Domain Objectives

- **8.1** Understand and integrate security in the Software Development Life Cycle (SDLC).
- **8.2** Identify and apply security controls in software development ecosystems.
- **8.3** Assess the effectiveness of software security.
- **8.4** Assess security impact of acquired software.
- **8.5** Define and apply secure coding guidelines and standards.

Domain 8 focuses on the software security piece of the CISSP exam. In this domain we will discuss how security is integrated into the software development life cycle, which is a formalized framework used to guide software developers in building code that is not only functional but secure. We will also examine the how security controls are applied in software development environments; how to determine the effectiveness of software security controls; and how to assess the impact on security of software that is acquired from a variety of sources, including software that is bought, commissioned, or even developed internally. We will also explore some key concepts regarding secure coding guidelines and standards.



Objective 8.1 Understand and integrate security in the Software Development Life Cycle (SDLC)

In this objective we delve into the basic foundations of the software development life cycle. We will discuss several aspects of this framework, including development methodologies and maturity models. We will also describe how the functions of operations and management, as well as change management, and the integrated product team contribute to integrating security into the SDLC.

Software Development Life Cycle

The software development life cycle (SDLC), as mentioned earlier, is a formalized process used to provide structure and guidance to software development efforts in an organization. The term *life cycle* implies, naturally, that there is a beginning and end to the life of a software application. The SDLC begins with the software's creation and ends with its retirement or disposal. Several different versions of the SDLC are in use and they vary depending on the industry segment and specific requirements levied by customers, professional bodies, governmental agencies, and even the organization itself. We'll discuss a few of the key SDLC methodologies in this objective.

Generically, the software development life cycle consists of the following phases, during which specific activities occur:

1. **Requirements gathering** During this phase, requirements are formally gathered, including functional requirements (*what* the software must do) and performance requirements (*how well* it must do it).
2. **Design** This phase involves determining *how* the software will meet its functional and performance requirements.
3. **Development/acquisition** During this phase, software is either programmed internally or acquired from a third party, meaning it is commissioned or bought from another source.

4. **Testing** This phase involves verifying that the specific functional and performance requirements are implemented and working as expected; this phase also involves integrating software code with other applications and systems to ensure they are interoperable.
5. **Operations/maintenance** This phase involves implementing software after it has been accepted by the organization and putting it into operational use. Maintenance involves planned change management, updates, and so on.
6. **Retirement/replacement** The software is eventually replaced with a newer version, retired, or otherwise decommissioned.

Note that not every phase mentioned here is included in every SDLC model; in fact, other SDLC models may include additional phases. Understand that these are only generic phases presented for familiarization purposes to meet the exam objectives. We will discuss specific development methodologies in the next few sections.



EXAM TIP While various SDLC models may contain different phases or use different names for the phases, you should be familiar with the generic SDLC model and its phases for the CISSP exam.

Development Methodologies

Where an SDLC model gives you an overall framework that delineates the process of developing and fielding software applications in "phases," development methodologies are more detailed and provide specific processes framed around an SDLC. They may also appear as SDLCs in their own right. Most earlier methodologies did not include security as a process or a phase. However, some of these methodologies have since been modified to build security into the development process, versus attempting to "bolt on" security as an afterthought once the software has already been built or even deployed.

We are going to discuss some of the traditional development methodologies in this section, but keep in mind that they have evolved over the years. Also, understand that the following discussions of the methodologies describe how they work in theory; you will rarely find them implemented in their "pure" form in the real world. Many of these methodologies have been modified, combined, or integrated with other methodologies to suit the needs of the developing organization.

Waterfall

The Waterfall methodology is the traditional software development method that harkens back to the very beginning days of software; for several decades it was the de facto development methodology of choice for the U.S. government. The Waterfall methodology is extremely structured and sequential and requires coordinated reviews, approvals, and other gatekeeping

events to move from one phase to the next. Because it is typically a rigid, linear, one-way methodology, it impedes developers' ability to go back to an earlier phase for rework if something goes wrong later in the development process.

Incremental and Spiral Methodologies

The Incremental methodology allows a software development team to produce software code that progresses a bit faster throughout the development phases (e.g., requirements gathering, design, development, testing, operation and maintenance, and retirement/replacement), albeit without all the necessary requirements that the final product should meet. Each iteration that completes the development life cycle is a finished, workable product, but may not include all the functions the software must eventually contain. While one piece (or functional area) of the software is iterating through the development process, another portion has already entered the process as well, so eventually all of the functional requirements the software must contain are added incrementally through several iterations. In fact, this methodology is sometimes called the *iterative methodology*. You could think of this methodology as multiple instances of the Waterfall methodology combined together.

The Spiral developmental methodology is a more formalized version of an iterative methodology. Each iteration, or *spiral*, is carefully planned, so the end user understands exactly which requirements each iteration will fill, and how the next spiral will be improved. Each new spiral has specific objectives it must meet. In addition to iterative development, each new version of the software is also thoroughly tested and analyzed for business risk throughout the development process.

Prototyping

Prototyping is a methodology that uses software models or prototypes that can be developed and analyzed to make sure they are meeting customer requirements. The customer may frequently receive a prototype version of the software for testing and approval; the danger of this is that the customer will get used to what the prototype looks like instead of the finished product or begin using the prototype in actual production. In some cases this is actually planned for and the prototype that is delivered to the customer is a functional, usable version of the software (called an *operational prototype*), albeit without all final requirements being met.

Another version of this methodology, called *rapid prototyping*, is exactly what its name describes. Prototypes are quickly developed to give to the customer for testing and use, often to fill an urgent business need and to get the customer a piece of code that fills a functional requirement. Although prototypes are generally discarded after development and testing, having fulfilled their function, *evolutionary prototypes* are built and presented to the customer with the intention of incremental improvement. In this approach, the prototype is given to the customer with the core functionality base, which is added to and modified over an incremental development cycle. Each prototype receives customer feedback that is incorporated into improvements for the next version of the prototype.

Rapid Application Development

Rapid Application Development (RAD) can be seen as an extension of prototyping; the major difference, however, is that less effort goes into upfront planning in the RAD environment. Planning and development are integrated, producing rapid prototypes the customer can review, provide feedback on, and either use or discard. There is less focus on the deliverable piece of software code for a particular iteration that meets all the functional and performance requirements laid out in a requirements gathering phase; the requirements are discovered and built-in as the application is developed and tested by the customer.

RAD evolved into a serious development methodology because other, more structured methodologies, such as Waterfall, Incremental, and Spiral, were too slow to keep up with the rapidly changing environment of customers. Often, by the time the software was finally developed to meet the original requirements, those requirements had changed or become obsolete due to the operating environment, technology, market, and so on. The RAD theory is that a 60 percent solution right now is better than a 100 percent solution in three years, with the understanding that over the course of iterative development, the solution would approach the final needs of the customers faster.

Agile

Agile is more of an approach to software development than a single methodology; in fact, many methodologies follow the Agile approach. As its name indicates, Agile promotes agility in software development that is not restricted by the overuse of structured processes that can get in the way of actual development. At the same time, Agile methodologies typically provide more structure than prototyping and RAD methodologies, but the structure is geared to assist in understanding the customer's requirements versus achieving software perfection over a longer period of time. A strength of the Agile approach is that it draws from all other methodologies and uses the best techniques in each. Agile can be incremental and iterative, but also flexible enough to adapt to the needs of each application. Agile methodologies focus on delivering small increments of functional code that are based on business need, versus delivering one large, monolithic application that may take years to perfect.

A key characteristic of Agile methodologies is a *user story*. This essentially is a use case from the user's point of view, which states exactly what a user wants to do (function), why the user wants to do it (result), and how (performance). This user story is effective in telling the developer exactly what the end customer wants to see in the software functionality. This approach may be more effective than formal requirements development because it allows the user to express what they want in their own words.

Key Agile methodologies include

- **Scrum** This allows for development points to be reset as product features need to be added, changed, or removed; it uses a two-week development interval called a *sprint*, after which new features or changes are delivered.

- **Extreme Programing (XP)** This approach is similar to Scrum, except without the sprints and with more code review using an approach called *pair programming*, where programmers work in pairs and are constantly inputting and checking each other's work. XP also is unique in that developers write test cases before the code is even written, so the code is constantly improved until it actually succeeds at passing the test case.
- **Kanban** Kanban was originally a production scheduling system developed by Toyota, but eventually software developers adapted it. It uses visual tracking of tasks so all team members understand what the priorities are and what they should be working on at that particular time.

DevOps and DevSecOps

In traditional development environments, software developers and operations personnel were two separate groups, with little productive interaction between them. This often resulted in software that did not meet its customers' requirements, was faulty, or became subject to rework. Over time, a new paradigm of combining these two functional entities during software development emerged, known as DevOps. Now integrated development teams consist of developers and operations personnel. These two broad areas include programmers, software engineers, IT personnel, and quality control team members. The benefit of this new paradigm is that developers and operators can work together to produce code that better meets more realistic customer requirements and can be adapted as needed when those requirements change.

Although DevOps is an advancement in integrating direct functional areas for producing better-performing code, it leaves out the security piece of the equation. Code still tends to be developed in a security vacuum, such that security considerations are not included in the development process. Although DevOps makes code better from a functionality perspective, it leaves it very vulnerable to increasingly sophisticated software attacks. A newer paradigm, known as DevSecOps, integrates security personnel into the DevOps team, allowing security consideration to be included in the development process, in the application code itself from the beginning of development, rather than attempting to add it on to a complete piece of software after the fact.



NOTE DevOps and DevSecOps really aren't development methodologies per se, they are simply a method of combining personnel from different functional areas to contribute to the development process. Both can use a variety of development methodologies within the team.

Maturity Models

A maturity model is not a software development methodology. It is a formalized model for determining how well organized, defined, and managed an organization's development effort is. Maturity models provide a measurement or view of an organization's software development

capabilities, classifying them as loose and ad hoc on one end of the spectrum, or highly organized, defined, and precise on the opposite end. Most organizations fall somewhere in between these two extremes. The purpose of a maturity model is to give an organization a method and framework to help improve its software engineering and development capabilities. Although there are many different maturity models that span different functional areas, including security, software development, engineering, and so on, we will discuss in the following sections two key maturity models that you will need to know for the CISSP exam.

Capability Maturity Model

A Capability Maturity Model (CMM) provides a method to evaluate an organization's software development program along a defined spectrum of increasing effectiveness and identify areas for improvement. The maturity of an organization's software development effort can be gauged by inspecting the quality of its management, software development personnel, training, development methodology, and, in the end, the quality of its software.

Capability Maturity Model Integration (CMMI), originally developed by Carnegie Mellon University, and now managed by ISACA, is a popular maturity model that provides measurements of the software development program in an organization. It also provides guidance on the different phases of a software development life cycle and identifies best practices organizations should be using in their software development efforts. CMMI consists of six maturity levels:

0. **Incomplete** This level is considered ad hoc with no known formal processes.
1. **Initial** At the *initial* level, the development effort is unpredictable and reactive to circumstances; the development effort is often over budget and behind schedule.
2. **Managed** At the *managed* maturity level, there are the beginnings of a management infrastructure in place to control the software development effort at the project level, using basic planning techniques and metrics.
3. **Defined** At the *defined* level, there are formal organization standards in place that provide structure for the software development effort.
4. **Quantitatively Managed** At the *quantitatively managed* level, the organization has formal processes to not only manage the software development effort, but to extrapolate and analyze data to gain efficiencies and a roadmap for more effective management.
5. **Optimizing** At the *optimizing* level, the organization focuses on continuous process improvement and long-term software development program strategy.



EXAM TIP You should be familiar with the six levels of the CMMI model for the exam.

TABLE 8.1-1 Software Assurance Maturity Model

Governance	Design	Implementation	Verification	Operations
Strategy and Metrics	Threat Assessment	Secure Build	Architecture Assessment	Incident Management
Policy and Compliance	Security Requirements	Secure Deployment	Requirements-driven Testing	Environment Management
Education and Guidance	Security Architecture	Defect Management	Security Testing	Operational Management

Software Assurance Maturity Model

The Open Web Application Security Project (OWASP) is a professional organization whose stated goal is to improve the security of software. Toward that goal, OWASP has developed its own Software Assurance Maturity Model (SAMM), which is specifically focused on secure software development. Any organization can use the model, regardless of the size of their development effort. SAMM is broken into five critical business functions. The idea behind the model is that an organization that produces software should perform these five functions. Each of these functions is further broken down into three security practices, which contain security-related activities within each function. The total of 15 security practices can be assessed for effectiveness and maturity, which in turn gives the organization an idea of the maturity level at which it currently resides, and which maturity level it needs to pursue.

The five business functions that comprise SAMM, and their respective security practices, are summarized in Table 8.1-1.

Note that each security practice also consists of multiple activities decomposed down to additional levels.

Operation and Maintenance

Regardless of the development methodology or life cycle model that is used, once the software has been developed, accepted by the user, and implemented, it goes into operation and maintenance (O&M). However, during O&M, development does not necessarily stop. Often patches or upgrades are produced, based upon new functional or performance requirements. Security mechanisms may be introduced into the software, and previously undiscovered flaws may be identified and scheduled for mitigation during another update or iteration of the software. Version control is of critical importance after implementation to ensure that only authorized production code is in the production environment.

In addition to new requirements, updates, flaw remediation, and other expected changes to software, the vulnerability management process can affect software during O&M. New vulnerabilities are discovered every day, either in applications, their code libraries, or even in components that are incidental to the software, such as network connectivity modules or authentication mechanisms. These vulnerabilities must be addressed through patching, secure configuration, and overall security management.

Change Management

Change management is also a critical consideration in software development. No application ever survives its entire life cycle without changing to some degree, whether those changes are minor, such as a patch or an update, or even a major redesign of the software itself. Change management must be a formalized process with structure, rather than an ad hoc effort. Changes must always be carefully considered, tested, and documented.

While the formal change management process is necessarily an overarching management process, *change control* happens at a lower level, and is the process of controlling specific changes in a software application during its life cycle, and documenting those changes. Note that this can also be referred to as *software version control*. This is very similar to the configuration management process for a system in that it happens at the system or software level. The formal change and control management process essentially consists of requesting the change, testing the change, approving the change, implementing the change, and documenting the change; these formal processes may vary from organization to organization, depending on their software development maturity level.

Cross-Reference

The details of the formal change management process were discussed in Objective 7.9.

Integrated Product Team

An *integrated product team (IPT)* consists of team members with expertise in a wide variety of areas. This enables programmers to work with all manner of operations personnel, including IT, accounting, marketing, engineering, and other functional areas not traditionally involved with software development. This gives the developers insight into how these areas actually work, allowing them to produce better code that actually meets the needs of their customers. Note that an IPT is not necessarily a development methodology, but more of a management technique that contributes to software design and security by allowing the right people from a variety of functional areas to be involved in the software development process. However, IPTs often use Agile development methodologies.

REVIEW

Objective 8.1: Understand and integrate security in the Software Development Life Cycle (SDLC) In this objective we began our discussion of software development security. We discussed the importance of the software development life cycle in providing a structured framework upon which to build functional code and applications. We explored several different development methodologies, which can be independent of the SDLC. These include Waterfall, which is a rigid, sequential methodology, Incremental and Spiral (both of which allow multiple iterations of a development phase), prototyping, Rapid Application Development, and Agile. Agile is more of a development approach than a methodology and includes several methodologies such as Scrum, Extreme Programming, and Kanban.

We also discussed both DevOps and DevSecOps, which resulted from the need to combine developers with operations and security functions to produce software that is closer to the actual user requirements, is functional, and secure.

We also discussed maturity models, which allow an organization to gauge its level of maturity in its software development processes. CMMI and SAMM are the two specific maturity models we discussed. We next briefly explored the need to ensure software is kept up-to-date and patched during operation and maintenance. Of particular importance for software O&M is vulnerability management, understanding that even after software implementation, security flaws will be discovered that must be remediated.

We also briefly mentioned the change management process, also discussed in Objective 7.9, and how it applies to software development. Finally, we discussed the value of an integrated product team, which not only combines developers with operators and security personnel, but also pulls in personnel from a variety of functional areas, including accounting, marketing, engineering, and other operational areas.

8.1 QUESTIONS

1. You are working on a contract with a government agency and are reviewing documentation for a legacy software application that must be completely rebuilt from scratch. The documentation indicates that there were key decision points during development of the software, after which required updates and changes were not allowed to be made. This likely affected the features and functions of the software. Which of the following software development methodologies most likely was used in developing this legacy application?
   - **A.** Incremental
   - **B.** Waterfall
   - **C.** Agile
   - **D.** Rapid Application Development
2. Which of the following development methodologies allows developers to build a quick model of software and provide it to the customer for review?
   - **A.** Prototyping
   - **B.** Waterfall
   - **C.** Spiral
   - **D.** Incremental
3. Which of the following best describes a relatively new software development methodology that integrates developers, operations (i.e., end users), and security personnel into the same team to develop software that is closer to a customer's requirements, is functional, and has security built-in?
   - **A.** Extreme Programming
   - **B.** Scrum

C. DevOps

D. DevSecOps

4. Which of the following maturity models is composed of business functions and corresponding security practices?

A. Scrum

B. CMMI

C. SAMM

D. Agile

8.1 ANSWERS

1. **B** This describes the Waterfall methodology of software development, since after decisions are made, they can rarely be reversed, and development cannot fall back to a previous phase for updates or new features.
2. **A** Prototyping allows a software developer to develop a model of the software and provide it to the customer for review and feedback. The prototype may or may not have all of the required functionality.
3. **D** DevSecOps combines representatives from the developer, security, and operations functional areas to jointly develop software that more closely meets customer requirements and is secure.
4. **C** The OWASP Software Assurance Maturity Model (SAMM) consists of five critical business functions, including Governance, Design, Implementation, Verification, and Operations, each of which is further composed of three security-related practices.



Objective 8.2 Identify and apply security controls in software development ecosystems

In this objective we will examine some of the security controls you should implement during the software development process and in the development environment.

Security Controls in Software Development

For much of this book we have discussed security controls, in the context of identification, authentication, authorization, auditing, accountability, nonrepudiation, cryptography, and physical security. Implementing security controls in software development ecosystems is equally critical. For the most part, administrative and technical controls are implemented

during the software development life cycle. These controls include software development policies, standardized secure development methodologies, and access control to production code, among others; but physical controls, such as physical protections for software repositories and development laboratories, are also necessary.

In this objective we will discuss a variety of security controls that an organization should implement in its software development environment, but we'll also cover foundational areas of software development that contribute to security, such as programming languages, software delivery, and software configuration management. We'll also discuss important topics such as code repositories, application security testing, and the orchestration and automation of security tools within the organization.

Programming Languages

All software is written as code, which is a set of instructions that tells the application how to execute. Software code is written using a programming language, which uses specific syntax, statements, variables, functions, subroutines, and so on. Programming languages have significantly evolved since the beginning of software development, from the use of arcane, complex hexadecimal instructions to structured, procedural code, to object-oriented programming, and finally, to natural language. There are five defined generations of programming languages that you should remember for the CISSP exam. Table 8.2-1 lists and briefly describes these five generations of languages.

TABLE 8.2-1 The Five Generations of Programming Languages

Generation	Level	Description
First	Machine language	Language in a format that the computer's processor can understand and work with directly
Second	Assembly language	Hardware-specific symbolic representation of machine language instructions
Third	High-level language	Structured programming languages that use defined statements and abstract instructions to the hardware level; examples include C++, Python, Fortran, and COBOL.
Fourth	Very high-level language	Domain-specific languages such as database management, web development, etc.; examples include SQL and HTML
Fifth	Natural language	Introduces knowledge-based programming, artificial intelligence (AI), and natural language constructs

Note that as each generation of languages progressed, two key changes occurred:

- Languages became increasingly easier for programmers to write in, providing much more functionality while requiring far less programming. For example, writing instructions for a task in assembly language would have taken considerably more programming effort than writing the same task in a fourth- or fifth-generation language.
- The level of abstraction increased. As languages progressed, the level of detail the programmer needed to understand about the underlying operating system and hardware *decreased*. This made programming easier, more efficient, and actually far more portable, enabling cross-platform programming, since the increase in abstraction enabled software to communicate with a wider variety of hardware platforms.

Regardless of the generation of programming languages, instructions eventually have to be converted or translated down to the machine level. Even a fifth-generation language must go through several layers of translation before a specific hardware platform or processor can understand its instructions. Software code can be translated using three different methods:

- **Assembler** Converts assembly language source code into machine code.
- **Compiler** Converts high-level language statements into machine-level formats specific to the underlying hardware and processor.
- **Interpreter** Interprets individual program statements, such as those found in scripts, directly from a higher-level code into machine code. Interpreted code is much more portable and does not have the larger overhead of being compiled into an executable.

A *code module* is a name for a defined piece of code, such as a subroutine, function, or other segment of software code that can be delineated from others. Usually a code module is characterized by its ability to perform a specific task or function in the application. There are two key characteristics of code modules you should understand for the CISSP exam which relate to program efficiency and, in turn, security. *Cohesion* refers to how many different types of tasks a code module is capable of carrying out. The fewer tasks that a module has been written to perform, the higher the level of cohesion. The higher the cohesion, the better.

Coupling refers to the interaction or communication that a code module must have in order to carry out its tasks. The lower the coupling, the less the module needs to interact or communicate with other modules to carry out its tasks. The higher the coupling, the more communication needed. Low (or *loose*) coupling is more desirable than high (or *tight*) coupling, as a matter of programming efficiency. Note that both high cohesion and low coupling are classic examples of the security concept of lessening the attack surface for the code.

Libraries

A software library is a collection of software components, such as pieces of code, functions, and so on, that can be used across many different software development projects. Software libraries are essentially collections of reusable code. Someone develops a piece of code that performs a needed function and saves this in the software library. Examples of these functions that might

be used to implement security controls, and include cryptographic modules, authentication mechanisms, and so on. These are the types of software libraries that can come in handy for another application if they are found to be solid and useful. Software libraries are often represented as code modules, executables, dynamic link libraries, and so on. Developers need to understand how the software code module is constructed and may need to develop an application programming interface (API) to integrate the module into their application if it does not easily integrate into the application.

Tool Sets

A tool set used by a development team may not be standardized between organizations, or even within the same organization. Every programmer has their favorite tool set, but an effort must be made to standardize these tool sets within the organization so that standardized development methodologies and predictable results are possible. A developer's tool set can come in many forms: customized scripts, homegrown applications, databases, spreadsheets, custom libraries, and many other nonstandard tools implemented to make a developer's job easier. Standardizing a developer's tool set can be difficult, but if management listens to the developers and adopts the best-of-breed tools that developers really need to do their jobs, along with incorporating a standardized integrated development environment (discussed next), it is possible. Note that these tool sets, in addition to being used for general software programming and development tasks, also typically include security functions, such as encryption and authentication mechanisms.

Integrated Development Environment

An integrated development environment (IDE) is a suite of tools designed for code development and leveraged through a single graphical user interface. In the early days of software development, code was written strictly as sequential lines of text. Today's IDEs allow object-oriented programming (OOP) using a visual graphical interface. IDEs support much more than simply writing lines of code by enabling developers to frequently check code for errors, execute it in a controlled environment, and test it using various built-in or customizable test cases. IDEs can be used for a variety of languages, provided the IDE supports them. Sophisticated IDEs can also be used to control access to code repositories, implement version control, and log and audit changes to code. IDEs can also include security functions and controls that can be easily implemented in the software they are used to develop.

Runtime

Runtime is a name given to an environment in which code written for any platform can be executed. A runtime environment acts as a sort of miniature operating system or virtual machine container where the application, often called *portable* or *mobile code*, can run as it is natively written. You'll most often see runtime environments used with mobile code written for web-based applications, such as Java, for example. The runtime environment serves as a

layer between the code and the operating system that the code can communicate with, and in turn, translates the code's request for resources and the responses returned from the host. Because it serves as a layer between the operating system and the code that runs within it, a runtime environment can protect both the running code from non-secure conditions in its environment, and vice versa.

Continuous Integration and Continuous Delivery

The manner in which code is released into the production environment has become extremely critical, especially as development methodologies have changed and code is delivered to the customer more efficiently. You should understand two key delivery concepts for the CISSP exam, continuous integration and continuous delivery (CI/CD).

Continuous integration means that as soon as the code is developed, tested, and approved, it is integrated into the test (or pre-staged) environment. This is very much an automated process; code is developed and merged into the code repository, where a series of automated tests take place. If the tests yield good results, the automated process delivers the code into the repository and a new software version is produced. Errors are identified quickly and can be corrected much more efficiently.

Continuous delivery allows for a constant, incremental rebuild of software versions that can be released at any time. Again, this is far more of an automated process than traditional delivery methods, so code can be released into production almost in real time, provided it meets all the testing and acceptance requirements.

Because both of these processes rely heavily on automation, it's important that code repositories and the automation pieces are kept secure. A popular CI/CD platform known as *Codecov* was compromised in 2021, allowing malicious actors to potentially infiltrate the code bases of any product or organization that used the platform.

Security Orchestration, Automation, and Response

Security orchestration, automation, and response (SOAR) is a new way of managing the wide assortment of security tools, such as suites, scripts, and executables, that organizations use to manage security processes within the organization. Typically implemented as an integrated platform, SOAR attempts to bring all of the various tools under one umbrella to make them easier to manage and integrate. SOAR supports common data formats so that security information can be easily exchanged between tools; in some cases APIs are available or developed to connect one application to another.

Enabling an organization to execute the functionality of each disparate tool and receive the results through a single user interface saves time and creates efficiencies in security management. SOAR can accelerate security responses by automating standard remediations, allowing analysts to focus on higher-priority events that require human analysis. Many security information and event management (SIEM) suites can also function as SOAR platforms, although additional add-on modules may be required, or the organization may custom-build interfaces between some of its homegrown tools and the SOAR platform.

Software Configuration Management

As its name indicates, software configuration management (SCM) is a management process. It is the process used to ensure that software is built and configured to an exacting set of standards, that it is documented, and that its versions are strictly controlled. Much of this is done through automated processes, although manual tasks and human management must be present to ensure that these automated processes work to maintain a controlled code base. Configuration management and change control were discussed in detail in Objective 7.9 and briefly in Objective 8.1, and the principles discussed in those objectives also apply to software configuration management.

Documentation and version control are likely the two most important aspects of software configuration management. Software must be documented in several respects. The code itself must be commented and thoroughly documented so that programmers who review it in the future will understand how and why the code was designed. Higher-level manuals detailing how to operate an application must also be written; sometimes security mechanisms must be documented for users or technicians.

Version control means that strict adherence to change and configuration control processes must be followed. Any change to the official code base must be documented and version numbers must be incremented, whether they are major or minor version numbers. The code base versions that are put into production must be separated from the code versions that are in development and test repositories. *Code synchronization*, which ensures that all working versions are synchronized with the latest approved changes, can be an issue if developers are allowed to randomly work on code without going through a formal code checkout process. This process, discussed in the next section, ensures that code synchronization issues are minimized.

Code Repositories

A code repository can be both a physical and logical secure storage area for code management. Code repositories function as a code library, where the various versions of software are stored and managed. Access control is critical for code repositories; developers should have access to development code only, while IT personnel and others responsible for actually implementing code into production environments should only have access to production code. Since versioning is a critical function of software configuration management, the ability to audit and identify changes to code must be part of any code repository system. Code in various stages, such as development, test, and production, should be "checked out" by an individual so that auditing can be implemented; changes can be locked or allowed, but attempts should be logged and audited.

Application Security Testing

As part of the majority of SDLC models, application security testing takes place after development but before software is implemented and goes into operation and maintenance.

There are several different types of testing that can take place, in addition to several testing methods. The typical types of application security testing that you will see during the SDLC include the following:

- **Unit** Tests individual pieces of code for performance and functionality.
- **Interface** Tests the interface between the code and other components of the application, as well as hardware.
- **Integration** Tests that the integration of the software application code with other system components does not reduce functionality or security in those components; this is closely related to interoperability testing as well.
- **Use case** Puts code through the paces for the various use cases that have been developed from user requirements, which ensures that the code performs its intended functionality under predefined conditions.
- **Security** Assesses the security of the software before implementation; includes vulnerability testing, penetration testing, and misuse case testing (also discussed in Objective 6.2).

Application code is tested for a variety of potential issues. These include ordinary programming flaws, known programming language vulnerabilities (such as certain unsecure functions in C++), and, of course, security issues. The following list recaps the security aspects of code testing mentioned in Objective 6.2:

- Input validation
- Secure data storage
- Encryption and authentication mechanisms
- Secure transmission
- Reliance on unsecured or unknown resources, such as library files
- Interaction with system resources, such as memory and CPU
- Bounds checking
- Error conditions resulting in a nonsecure application or system state

Cross-Reference

The various types of tests, as well as some of the security issues they test for, were also discussed in Objective 6.2.

The CISSP exam requires you to know not only the types of application security tests and the issues tested for, but also a few of the key test methods that are used. These include both static and dynamic application security testing, as well as others, discussed next.

Static Application Security Testing

Static application security testing (SAST) involves testing application code without actually executing it. SAST requires access to the actual source code rather than the compiled executable. SAST can be performed in several ways, including through manual or automated means. With manual SAST, the code is reviewed line-by-line by another programmer; this is a basic *code review* technique. Automated testing uses software suites that are specifically designed to rapidly go line-by-line through the application code to look for syntax errors, known flaws, faulty logic or sequencing, improper variable use, and many other common issues. Some of the more sophisticated software suites can look for issues such as lack of input validation, faulty function calls, and so on. Additionally, customized tests can be added to the software suites based on the needs of the organization.

Dynamic Application Security Testing

Dynamic application security testing (DAST) is performed while the application is in production mode. Once static testing takes place, the application is tested while actually executing. Since DAST is run during code execution, it does not require access to the source code. Specific software suites used for DAST can identify a variety of issues, including hardware resource overuse, lack of interoperability with other applications or operating systems, and so on. The test suite can also create artificial error conditions to see how the software code reacts.



EXAM TIP SAST requires access to the source code and is essentially a code review using manual or automated means. DAST does not require access to the source code and tests the code during execution.

Forms of Application Security Testing

Whether using SAST or DAST, application security testing has a specific goal, like unit testing or interface testing. Application security testing can take several forms, including:

- Vulnerability testing, during which known vulnerabilities in software code are discovered.
- Penetration testing, which attempts to "hack" the code using a variety of manual and automated methods, in an effort to circumvent security mechanisms. Ideally, the results of vulnerability and penetration testing are used to improve security mechanisms in the software code after the test.
- Misuse cases testing, which attempts to determine all the different ways software can be abused or misused. Misuse case testing is not always limited to security functions; it can also be used to check if users are using the software incorrectly for the functions it was intended. However, misuse often results in a security problem, such as an error (also called an *exception*) that may allow access to administrative privileges or the underlying operating system or hardware.

- *Fuzzing* is a form of misuse case testing that can be employed during security tests using both automated and manual means. Fuzzing takes nonstandard input and feeds it into the application while it is running to see how the application reacts. It essentially tests input validation, as well as variable usage and memory management.

REVIEW

Objective 8.2: Identify and apply security controls in software development ecosystems In this objective we discussed some of the security controls you should consider having in place during the software development effort. Some of these controls are specifically applied to the software development environment, such as access controls, but some of the characteristics of software development itself can help secure the development environment and the code, such as the type of programming language used, secure development methodologies, and so on.

We briefly covered some of the foundational aspects of software development, including a quick review of programming languages, libraries, tool sets, IDEs, and runtime environments.

We also discussed the concepts of continuous integration and continuous delivery, which enable errors to be quickly discovered and quality code to be put into production much faster. Security orchestration, automation, and response (SOAR) is an overarching management method used to integrate the various disparate security tools, such as applications, scripts, and so on, so that they can be managed under one interface and their results or outputs can be more efficiently collected and analyzed as usable information.

Software configuration management has two key characteristics: thorough documentation and version control for all software code. Code repositories are where the different official versions of code are stored; they must be logically and physically protected against unauthorized access and modification.

Application security testing is a key part of the SDLC. Applications can be tested using several methods, including static, dynamic, automated, and manual methods.

8.2 QUESTIONS

1. Emmy is a new programmer at your company. She has been tasked with updating a critical line-of-business application on a tight schedule. The current version of the application is fairly stable and leverages many standardized code modules used throughout other applications in the organization. If Emmy uses these code modules, which of the following is she taking advantage of?
   - **A.** Code libraries
   - **B.** Integrated development environment
   - **C.** Runtime environment
   - **D.** Continuous integration and continuous delivery

2. Caleb is a cybersecurity analyst who must perform security testing on a new application. He specifically needs to test input validation for the new application. Which of the following is the best type of testing to perform to meet this requirement?
   - **A.** Vulnerability assessment
   - **B.** Code review
   - **C.** Fuzzing
   - **D.** Unit testing

8.2 ANSWERS

1. **A** Reusing known stable code is taking advantage of the code libraries the organization maintains.
2. **C** Caleb must perform fuzzing on the application to test input validation, which involves inputting a variety of unexpected data into the application to see how the application handles those inputs.



Objective 8.3 Assess the effectiveness of software security

This objective presents methods for assessing the effectiveness of the overall security program with regards to software development. The same assessment principles we have addressed in other objectives throughout the book also apply to software security, such as testing, auditing, and risk management, and we will discuss these principles as they relate to software security.

Software Security Effectiveness

How does one assess the effectiveness of software security? Is it a measure of how many vulnerabilities are found in the software? Is the security effective only if the software never gets hacked? Is it effective because we use secure development methodologies or secure coding? These questions can only be answered by each organization, but there are several methods to assess if the software development process is secure, as well as assess the overall security of the software itself.

In both Objectives 6.2 and 8.2 we discussed the value of assessment and testing and the various ways to test software, such as static and dynamic testing. These methods tell us how secure software is before it is implemented. However, software must continue to be assessed after implementation, using those same methods, but also using auditing and risk analysis. We'll discuss those topics in the following sections.

Auditing and Logging Changes

Auditing and logging changes to software actually applies to two different contexts. First is the software development process. As mentioned in Objective 8.2, the official code base should be kept secure and should not be altered by unauthorized personnel. This is achieved by storing the code in a code repository and maintaining version control. There could be a version specifically approved for development, one for test, and, most importantly, one for production code. No one should have access to the code unless they are authorized, and even then, changes to any of those code bases should be approved and audited. Changes should be logged, particularly when changing production code, but also when changing development or test code. Version numbers should be incremented accordingly after changes are made, to maintain version control.

The other context for software auditing and logging involves the software itself, during operation and maintenance. Software should be developed with built-in logging and auditing mechanisms. For example, if someone attempts to change sensitive data in a software application, the application itself should log those attempts, or at least integrate with a larger logging system, such as the application log facility in Windows or a syslog server, to record the changes. Software also should log security events, such as authentication or authorization failures. Whether the software has its own built-in logging and auditing facilities or uses the facilities of the operating system or another logging application entirely, those logs should be protected from unauthorized personnel; only staff with administrative-level privileges should be allowed to access software application logs.



EXAM TIP Keep in mind the critical importance of logging all changes to a code base and auditing access to all versions of software code, to include development, test, and particularly production code.

Risk Analysis and Mitigation

General risk management principles were discussed at length in Objective 1.10. The same risk management principles apply to software, but in this objective, we will also present more specific details about how risk management can be used to analyze software security and mitigate security issues.

Applying risk management principles to the software development and security processes involves determining the threat actors, threats, and vulnerabilities associated with our software assets. For the purposes of this discussion, software assets include not only completed software applications, but all of the code bases in their specific versions: development, test, and production. Our software assets could also include the tools we use to develop and manage software, such as IDEs, code repositories, and various utilities. All these assets must also be secured against various threats and examined for vulnerabilities.

Vulnerabilities can be discovered in software using some of the techniques we discussed in Objective 8.2, such as SAST and DAST. However, another good place to start looking for vulnerabilities is to compare your software against industry vulnerability databases and lists, such as the Common Weakness Enumeration (CWE) Top 25 Most Dangerous Software Weaknesses, a community-developed list of hardware and software vulnerabilities hosted by MITRE, and the OWASP Top Ten, a list of the ten most critical vulnerabilities in web application software, published by the Open Web Application Security Project.

The OWASP Top Ten is periodically updated, so be sure to check for the latest version (https://owasp.org/Top10/). As an example, the OWASP Top Ten for 2021 is listed here:

- A01: Broken Access Control
- A02: Cryptographic Failures
- A03: Injection
- A04: Insecure Design
- A05: Security Misconfiguration
- A06: Vulnerable and Outdated Components
- A07: Identification and Authentication Failures
- A08: Software and Data Integrity Failures
- A09: Security Logging and Monitoring Failures
- A10: Server-Side Request Forgery (SSRF)

Once you have identified threats and vulnerabilities associated with your software assets, you can determine likelihood and impact, as you would with any other type of asset when assessing risk. The same types of qualitative and quantitative analyses introduced in Objective 1.10 apply to determine likelihood and impact with software assets.

Once you've ascertained software risks, the next step, of course, is to respond to the risks. As discussed in Objective 1.10, the four primary risk responses are mitigation, transfer or sharing, acceptance, and avoidance. Each of these can be employed in various ways for software risk. Mitigation involves patching or updating software with stronger security mechanisms; transfer or sharing risk may involve outsourcing some of the more complex coding efforts; acceptance may be an option after all other risk responses have been exhausted and the organization must accept some level of residual risk; and avoidance means that the organization would likely stop producing the software or follow another path in its software development efforts.

Cross-Reference

Risk management was discussed in great detail in Objective 1.10, and all the risk management principles it covered easily apply to software development efforts and software assets.

REVIEW

Objective 8.3: Assess the effectiveness of software security In this objective we reviewed some of the assessment methods discussed in various other objectives as they apply to the context of assuring software security effectiveness. In addition to software testing, assessing the effectiveness of software security requires software auditing and logging of changes to the code base. We also quickly summarized the risk analysis and response process as it applies to software security.

8.3 QUESTIONS

1. Which of the following is necessary to ensure version control and to detect any unauthorized changes in the code base?
   - **A.** Static code analysis
   - **B.** Dynamic code analysis
   - **C.** Auditing all changes
   - **D.** Manual code review
2. Your company is a software vendor that is performing a risk analysis of a software package before releasing it to the market. Your company has invested thousands of dollars and hours of labor to develop the software package. After intensive testing, several critical vulnerabilities are discovered in the software. Which of the following is the best risk response given that your company has invested so much time and money in the application?
   - **A.** Accept the risk and release the software.
   - **B.** Perform risk mitigation through patching, configuration changes, and updates.
   - **C.** Use risk avoidance by halting the software release completely.
   - **D.** Transfer the risk by purchasing cyber insurance on the software application.

8.3 ANSWERS

1. **C** Auditing all changes to the software code base is critical to maintaining version control and detecting any unauthorized changes to the code.
2. **B** Your company should perform risk mitigation to the extent possible through patching, configuration changes, and updates. Only after the residual risk is at an acceptable level should your company release the software.



Objective 8.4 Assess security impact of acquired software

Throughout the previous three Domain 8 objectives we have discussed security of software development within an organization. In this objective we continue the discussion of security in software but from the perspective of software acquired elsewhere, such as commercial software, open-source software, and third-party software created specifically for an organization. Each of these has its own unique challenges with software security.

Security Impact of Acquired Software

In this domain we've discussed ensuring that the software development life cycle is secure, as well as the software produced by an organization's own development efforts. In this objective we want to take a look at software acquired through other means. Not every organization develops its own software applications; in fact, for the most part, organizations acquire software as part of commercial-off-the-shelf packages, or, in the case of a specialized line-of-business software, possibly turnkey solutions or even software commissioned from third-party developers. The organization can also use software from the open-source community. We will look at security considerations for each of these software sources.



NOTE In all of these external sources of software, the same processes an organization would use to assure software security in its own development efforts can be used to establish trust in the software acquired from third parties. These include ensuring the third party uses secure software development methodologies, enforces strict access control to its code base, maintains secure code repositories, employs strict software versioning, implements security testing, releases patches and updates, and maintains overall secure management processes.

Commercial-off-the-Shelf Software

Commercial-off-the-shelf (COTS) software is normally produced by and purchased from companies like Microsoft, for example, or other mainstream software development companies. Think of office productivity suites, accounting software, video editing software, or music software. Typically, you don't actually buy software—you buy a license to use it, and the license comes with several different restrictions imposed by the developer. We won't go into the details of the many different licensing models available, but suffice it to say that the license for a piece of software determines how it is used, purchased, transferred, and so on.

For the most part, organizations get access to executables they can install on the system, but they don't get access to source code, since source code is the intellectual property the

software developer relies on to maintain its competitive edge in the marketplace. Most licenses prohibit disassembling or reverse engineering software, since that takes away the competitive edge of the software developer. Likewise, most licenses prohibit copying software or using it without purchasing it. Unlike the old days of simply inputting a software key, today's licensing mechanisms typically require Internet activation and periodic online updates to maintain the security of the licensing agreement.

COTS software may be the most secure source of software available, since it normally has been tested and implemented on a large scale. However, that doesn't mean COTS software is free of vulnerabilities. New vulnerabilities are discovered in commercial software all the time, and although most software development companies usually are able to release updates or patches to resolve a vulnerability or security issue before it is exploited, this isn't always the case. Zero-day vulnerabilities (vulnerabilities that can be exploited immediately and have no mitigations developed for them) have been discovered in most popular COTS software, prompting the software developers, anti-malware vendors, and security practitioners to scramble to quickly mitigate the flaws and prevent malicious entities from taking advantage and exploiting the software.

All software, whether an organization develops it internally or purchases or commissions it from another organization, should be scanned for vulnerabilities on a recurring, periodic basis. A well-managed patch management program should be in place to ensure that software is updated and patched on a regular basis.

Open-Source Software

While from a practical standpoint, most open-source software is benign, there is still the possibility that vulnerabilities exist in its code, even after long periods of time. The strength of open-source software is its community of developers, working diligently to ensure that (for the most part) open-source software that is released is stable and highly functional, in addition to being secure.

Of course, there is still the chance that open-source software can be infiltrated by malicious entities and compromised. The best defense for an organization that is contemplating using open-source software is to attempt to trace its origins back to verified reputable developers or communities. After an organization implements reputable open-source software, it should scan the software for vulnerabilities periodically, the same as for commercial software.

Third-Party Software

Third-party software is usually developed specifically for an organization that does not develop its own software. This could be a specialized piece of code developed for a line-of-business application, to run a production piece of machinery, or for a specialized business need. In any event, the key to assessing the security impact of third-party software is to assess whether the software developer is trustworthy. The organization must determine to what extent it can trust the third-party developer. It can do so by reviewing the third-party developer's reputation (via social media, websites, publicly available audit reports, and so on) and its software development certifications, such as CMMI level, for example.

After determining that a third-party software developer is trustworthy, the organization must determine how it will enforce security mechanisms in both the development and delivery of the software. This is where contracts, service level agreements (SLAs), and other legal mechanisms come into play. The organization has the ability, through contractual agreements, to enforce certain aspects of security, such as a verifiable secure software development methodology, code review and testing, required experience and skill sets of its programmers, and so on. With these things in place, the third-party developer is even more trustworthy.



EXAM TIP Third-party software development security can be achieved through trust in the third-party developer's processes, as well as contractual agreements between the organization and the third party.

Managed Services

Managed services present their own challenges, but in some ways, using managed services may be an easier way to manage the security of software, since most of the hard work is offloaded to the service provider. Managed services are services performed by a third party on behalf of an organization. These services range from managing the organization's infrastructure, servers, and security to managing its software, desktops, and other aspects of its information technology. All these services could also be managed as part of a cloud service offering, in the form of Software as a Service (SaaS), Infrastructure as a Service (IaaS), and Platform as a Service (PaaS). The provider is normally responsible for software security, including the development process and the security mechanisms that are part of the application. The provider typically is also responsible for keeping the software up to date with the latest patches and changes. The same applies if the provider is offering either IaaS or PaaS.

There are two keys to ensuring security for managed services, especially software services. First is the level of trustworthiness of the provider. You need to determine if the provider is reputable. You should perform due diligence by carefully researching the provider both online and through interviewing the provider to answer the following questions: Do they hire skilled and experienced developers? Do they use a sound software development methodology? Do they guarantee levels of service or security in their contract or agreement?

The agreement itself is the second key to ensuring security in managed services. Whether the agreement is a contract or SLA guaranteeing a specific level of security, it should protect the organization by specifying security responsibilities, levels of functionality and performance, and so on. The contract can be written to require the provider to use experienced developers for its software and to provide security audit reports or allow itself to be audited by the organization.

Cross-Reference

We discussed third-party service offerings, as well as cloud service providers, in Objectives 1.12, 3.5, and 7.7.

REVIEW

Objective 8.4: Assess security impact of acquired software This objective briefly discussed ways to assure the security of software that has been acquired through commercial purchase or license, open sources, third-party developers, or managed services. The same software assurance practices discussed in the previous Domain 8 objectives apply to these sources of software as well; the major difference is that the organization does not control the development or security processes. The two keys to assuring software security are trust in the developer and contractual requirements. Software security is assured also through some of the same processes that an organization would use if it had its own development efforts: testing, secure code repositories, software versioning, vulnerability testing, patches, and updates.

8.4 QUESTIONS

1. Your organization is looking to replace a software application for a specialized production system that has reached its end-of-life. The original manufacturer of the production system is out of business, so they are not available to help upgrade the software associated with the system. Your company has no internal development program. Which of the following would be the best solution for acquiring software to run the unique production system?
   - **A.** Commercial-off-the-shelf software
   - **B.** Third-party software developer
   - **C.** Open-source software
   - **D.** Internally developed software
2. Your company has contracted with a managed service provider to provide proprietary software services to support a specialized line-of-business application. The managed service provider owns the software licenses and takes care of its maintenance and patches. However, there is a new major version of the software that the managed service provider refuses to deploy for your company, stating that it is not contractually obligated to do so. The company requires this new version to accommodate upcoming changes to its production processes. Not upgrading the software could possibly cost the company a great deal of money, and its older software would also not be compatible with new security mechanisms coming as part of the process changes as well. If your company continues to require the software, which of the following should it consider regarding receiving future major versions and upgrades to the software?
   - **A.** Review statutes and regulations directing the protection of the data and software processes, and consider suing the provider in civil court.
   - **B.** Find a different third-party vendor whose software is compatible with your processes.
   - **C.** Amend the contract and service level agreement with the software provider to include a requirement for new versions of the software and security upgrades.
   - **D.** Make the decision to stop using the service provider and go with commercial-off-the-shelf software or open-source software.

8.4 ANSWERS

1. **B** Using a third-party software developer may be the best choice, since the organization does not have its own internal software development program. Since it is a unique production system, and the original manufacturer is out of business, COTS software is likely not a viable alternative, nor is open-source software, due to the unique, specialized nature of the production system.
2. **C** If the organization continues to require the software, it should amend the contract or service level agreement with the managed service provider to include major upgrades and new versions. Only after the organization determines the security and financial risk it would incur if it moves to a different software package should it reconsider the contract and go with COTS software or even open-source software.



Objective 8.5

Define and apply secure coding guidelines and standards

In this last objective of Domain 8, we're going to complete our discussion of software security by covering secure coding standards and guidelines, common software vulnerabilities, and software-defined security.

Secure Coding Guidelines and Standards

Writing code that is free of common security flaws is the key to developing secure software. Using even the most secure development environment makes no difference if the final software product is riddled with flaws and vulnerabilities. In the following section, we will discuss common security vulnerabilities that often make it into source code if secure programming practices are not followed. Next, we'll focus on the necessity of application programming interface (API) security. We'll then take a look at secure coding practices and why they are important. Finally, we'll briefly discuss the next logical extension of software-defined networking, which is software-defined security.

Security Weaknesses and Vulnerabilities at the Source-Code Level

Despite the best development efforts, even those that include a secure separate environment, code repository, and strong management practices, application source code will always have vulnerabilities. Sometimes these flaws are quite serious, and other times they are minor and only affect functionality. But even minor vulnerabilities can be chained together to lead to

a greater weakness. There are several common vulnerabilities that you can easily address in source code, but other vulnerabilities may require you to dig very hard to find them.

The best way to find vulnerabilities, of course, is through testing. We discussed both static application security testing (SAST) and dynamic application security testing (DAST) in Objective 8.2. Both types of tests look for common vulnerabilities that you can identify and remedy. Most of the vulnerabilities you will encounter are related to faulty input, resource misuse (e.g., memory, CPU, and storage space), vulnerable functions inherent to the programming language itself, and faulty security mechanisms (e.g., legacy or vulnerable authentication and encryption modules). Recall from Objective 8.3 that common vulnerabilities are listed in sources such as the OWASP Top Ten and the CWE Top 25 Most Dangerous Software Weaknesses. Table 8.5-1 shows the 2021 CWE Top 25 list as an example of the most common but critical vulnerabilities reported in 2021.



EXAM TIP You will not be expected to know the CWE Top 25 or the OWASP Top Ten verbatim for the exam, but you should still be familiar with some of the most critical software vulnerabilities and how secure coding practices can mitigate them.

Note that neither the OWASP Top Ten nor the CWE Top 25 is comprehensive; there are hundreds of other software vulnerabilities you should be aware of that are contained in different lists or databases. Many of these other resources focus on specific database, web application, or operating system vulnerabilities. There are also vulnerability lists specific to applications. Almost all of these vulnerabilities, however, can be mitigated through secure coding practices, discussed later in the objective.

Security of Application Programming Interfaces

An application programming interface (API) is a stand-alone piece of code, usually a module, that allows separate applications (and operating systems) to connect to each other and exchange data. The applications may have different data formats, security mechanisms (authentication and encryption, for instance), and methods of exchanging data. An API serves to bridge those gaps and allow seamless integration between the applications at their various levels, up to and including user interfaces. APIs can be constructed to be reusable, based on the different applications they connect to, and are often found in an organization's code libraries, enabling developers who are building a new application to incorporate an existing API to connect the new application to legacy applications, for example.

APIs have security issues that must be addressed as well. In addition to the normal common software vulnerabilities mentioned earlier, APIs must be programmed to accept a variety of inputs and translate them to data that is secure, formatted correctly, and useful to the other application. Parameter validation is important here; the API must determine if the data or commands it is receiving from the other application or OS are correct.

TABLE 8.5-1 CWE Top 25 (adapted from https://cwe.mitre.org/top25/archive/2021/2021_cwe_top25.html)

Ranking	Vulnerability
1	Out-of-bounds Write
2	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting")
3	Out-of-bounds Read
4	Improper Input Validation
5	Improper Neutralization of Special Elements used in an OS Command ("OS Command Injection")
6	Improper Neutralization of Special Elements used in an SQL Command ("SQL Injection")
7	Use After Free
8	Improper Limitation of a Pathname to a Restricted Directory ("Path Traversal")
9	Cross-Site Request Forgery (CSRF)
10	Unrestricted Upload of File with Dangerous Type
11	Missing Authentication for Critical Function
12	Integer Overflow or Wraparound
13	Deserialization of Untrusted Data
14	Improper Authentication
15	NULL Pointer Dereference
16	Use of Hard-coded Credentials
17	Improper Restriction of Operations within the Bounds of a Memory Buffer
18	Missing Authorization
19	Incorrect Default Permissions
20	Exposure of Sensitive Information to an Unauthorized Actor
21	Insufficiently Protected Credentials
22	Incorrect Permission Assignment for Critical Resource
23	Improper Restriction of XML External Entity Reference
24	Server-Side Request Forgery (SSRF)
25	Improper Neutralization of Special Elements used in a Command ("Command Injection")

Secure Coding Practices

There are many processes involved with secure coding. First, having a secure coding methodology that is standardized throughout the organization and required to be followed by programmers is a first step of the process. A second step is intense code review, often from peers, but also from formal review and test processes. Although secure development was traditionally

not included in most software development life cycle models, it has been increasingly integrated from the beginning of those life cycles, as early as the requirements stage using modern development methodologies, such as DevSecOps, for instance.

Cross-Reference

DevSecOps and other development paradigms were discussed in Objective 8.1.

Some of the more common secure software coding methodologies include the secure practices promulgated by OWASP, an independent professional organization that recommends software security methods (and publishes the previously introduced OWASP Top Ten). OWASP recommends establishing the following foundation for secure software development:

- Defined roles and responsibilities for security and development personnel
- Adequate security training for developers
- Implementation of a secure development life cycle model
- Adoption of secure coding methodologies

In addition to these overarching recommendations, OWASP also has published a Secure Coding Practices Checklist, which focuses on the following areas (adapted from *OWASP Secure Coding Practices Quick Reference Guide*, https://owasp.org/www-pdf-archive/OWASP_SCP_Quick_Reference_Guide_v2.pdf):

- Input validation
- Output encoding
- Authentication and password management
- Session management
- Access control
- Cryptographic practices
- Error handling and logging
- Data protection
- Communication security
- System configuration
- Database security
- File management
- Memory management

OWASP also includes in the checklist general coding practices, including some of the ones we've already discussed, such as the need for secure code review and testing, separate development environments, use of proven code modules and libraries, and so on.

A variety of other secure coding practices are available to the public, some developed and promulgated by independent professional organizations like OWASP, others by software or OS vendors, and still others by the U.S. government and industry standards bodies.

Software-Defined Security

Software-defined networks are faster, more efficient, and allow dynamic routing and switching, effectively replacing the slower traffic decisions made by hardware. Software-defined security takes it to the next level, enabling advanced traffic filtering to be performed as well. Software-defined security can perform both content-based filtering and context-based filtering, while filtering based on both simple and complex rule sets. This technology isn't just limited to traffic filtering, however; software-defined security can also perform the same functions as hardware firewalls, IDS, proxies, and other security devices. Software-defined security is also hardware-agnostic, independent of hardware vendor.

REVIEW

Objective 8.5: Define and apply secure coding guidelines and standards This objective completed our discussion of software development security (Domain 8). In this objective we discussed secure coding guidelines and standards. First, we talked about some of the common security weaknesses and vulnerabilities found in application source code. These can be identified during both manual and automated testing. Most of these security weaknesses and vulnerabilities are easily remedied through secure coding but may sometimes require occasional software redesign to accommodate security mechanisms. There are hundreds of software vulnerability sources you can find that will show critical vulnerabilities for databases, web applications, and so on, including the OWASP Top Ten and CWE Top 25 lists.

We also discussed the security of application programming interfaces. This is critical because the link between applications used to exchange data must be secure and provide for strong authentication and encryption mechanisms, as well as data integrity.

We also examined secure coding practices, which begin with management's requirement to implement standards in the organization. Secure coding practices include code review and extensive testing to locate and mitigate vulnerabilities. Secure coding practices and methodologies include those promulgated by independent professional organizations like OWASP, by software or OS vendors, and by the U.S. government and industry standards bodies.

Finally, we discussed software-defined security, an extension of software-defined networking. Software-defined security dynamically secures software-defined networks by making faster traffic-filtering decisions and fulfilling advanced security hardware functions.

8.5 QUESTIONS

1. Which of the following is critical to discovering vulnerabilities in code the organization has developed?
   - **A.** Secure coding standards
   - **B.** Testing
   - **C.** Secure design
   - **D.** Security requirements
2. You have been assigned to oversee security testing for a development effort focused on creating a new line-of-business application that must integrate with several existing applications. During testing, you discover that the authentication mechanisms between the new application and the existing ones automatically default to a legacy cryptographic algorithm. Which of the following should you address to ensure the authentication mechanisms between applications are secure by default?
   - **A.** Input validation
   - **B.** Network interfaces
   - **C.** Application programming interfaces
   - **D.** Memory and CPU usage

8.5 ANSWERS

1. **B** To discover vulnerabilities in software that the organization has developed, testing is absolutely necessary. The other choices are critical in preventing those vulnerabilities before and during development.
2. **C** Application programming interfaces are used to connect different applications together and may be responsible for exchanging authentication data. To ensure that the authentication mechanisms used between applications are functioning correctly, the APIs must be securely developed.

This page intentionally left blank



About the Online Content

This book comes complete with TotalTester Online customizable practice exam software with 300 practice exam questions.

System Requirements

The current and previous major versions of the following desktop browsers are recommended and supported: Chrome, Microsoft Edge, Firefox, and Safari. These browsers update frequently, and sometimes an update may cause compatibility issues with the TotalTester Online or other content hosted on the Training Hub. If you run into a problem using one of these browsers, please try using another until the problem is resolved.

Your Total Seminars Training Hub Account

To get access to the online content you will need to create an account on the Total Seminars Training Hub. Registration is free, and you will be able to track all your online content using your account. You may also opt in if you wish to receive marketing information from McGraw Hill or Total Seminars, but this is not required for you to gain access to the online content.

Privacy Notice

McGraw Hill values your privacy. Please be sure to read the Privacy Notice available during registration to see how the information you have provided will be used. You may view our Corporate Customer Privacy Policy by visiting the McGraw Hill Privacy Center. Visit the **mheducation.com** site and click **Privacy** at the bottom of the page.

Single User License Terms and Conditions

Online access to the digital content included with this book is governed by the McGraw Hill License Agreement outlined next. By using this digital content you agree to the terms of that license.

Access To register and activate your Total Seminars Training Hub account, simply follow these easy steps.

1. Go to this URL: **hub.totalsem.com/mheclaim**
2. To register and create a new Training Hub account, enter your e-mail address, name, and password on the **Register** tab. No further personal information (such as credit card number) is required to create an account.

   If you already have a Total Seminars Training Hub account, enter your e-mail address and password on the **Log in** tab.
3. Enter your Product Key: **pg5g-gpjg-js4v**
4. Click to accept the user license terms.
5. For new users, click the **Register and Claim** button to create your account. For existing users, click the **Log in and Claim** button.

   You will be taken to the Training Hub and have access to the content for this book.

Duration of License Access to your online content through the Total Seminars Training Hub will expire one year from the date the publisher declares the book out of print.

Your purchase of this McGraw Hill product, including its access code, through a retail store is subject to the refund policy of that store.

The Content is a copyrighted work of McGraw Hill, and McGraw Hill reserves all rights in and to the Content. The Work is © 2023 by McGraw Hill.

Restrictions on Transfer The user is receiving only a limited right to use the Content for the user's own internal and personal use, dependent on purchase and continued ownership of this book. The user may not reproduce, forward, modify, create derivative works based upon, transmit, distribute, disseminate, sell, publish, or sublicense the Content or in any way commingle the Content with other third-party content without McGraw Hill's consent.

Limited Warranty The McGraw Hill Content is provided on an "as is" basis. Neither McGraw Hill nor its licensors make any guarantees or warranties of any kind, either express or implied, including, but not limited to, implied warranties of merchantability or fitness for a particular purpose or use as to any McGraw Hill Content or the information therein or any warranties as to the accuracy, completeness, correctness, or results to be obtained from, accessing or using the McGraw Hill Content, or any material referenced in such Content or any information entered into licensee's product by users or other persons and/or any material available on or that can be accessed through the licensee's product (including via any hyperlink or otherwise) or as to non-infringement of third-party rights. Any warranties of any kind, whether express or implied, are disclaimed. Any material or data obtained through use of the McGraw Hill Content is at your own discretion and risk and user understands that it will be solely responsible for any resulting damage to its computer system or loss of data.

Neither McGraw Hill nor its licensors shall be liable to any subscriber or to any user or anyone else for any inaccuracy, delay, interruption in service, error or omission, regardless of cause, or for any damage resulting therefrom.

In no event will McGraw Hill or its licensors be liable for any indirect, special or consequential damages, including but not limited to, lost time, lost money, lost profits or good will, whether in contract, tort, strict liability or otherwise, and whether or not such damages are foreseen or unforeseen with respect to any use of the McGraw Hill Content.

TotalTester Online

TotalTester Online provides you with a simulation of the CISSP exam. Exams can be taken in Practice Mode or Exam Mode. Practice Mode provides an assistance window with hints, explanations of the correct and incorrect answers, and the option to check your answer as you take the test. Exam Mode provides a simulation of the actual exam. The number of questions, the types of questions, and the time allowed are intended to be an accurate representation of the exam environment. The option to customize your quiz allows you to create custom exams from selected domains, and you can further customize the number of questions and time allowed.

To take a test, follow the instructions provided in the previous section to register and activate your Total Seminars Training Hub account. When you register, you will be taken to the Total Seminars Training Hub. From the Training Hub Home page, select your certification from the Study drop-down menu at the top of the page to drill down to the TotalTester for your book. You can also scroll to it from the list of Your Topics on the Home page, and then click the TotalTester link to launch the TotalTester. Once you've launched your TotalTester, you can select the option to customize your quiz and begin testing yourself in Practice Mode or Exam Mode. All exams provide an overall grade and a grade broken down by domain.

Technical Support

For questions regarding the TotalTester or operation of the Training Hub, visit **www.totalsem.com** or e-mail **support@totalsem.com**.

For questions regarding book content, visit **www.mheducation.com/customerservice**.

This page intentionally left blank

Index

A

ABAC (attribute-based access control), 243
absorption in wireless technologies, 198
access audits for facilities, 385–386
access control
    authentication. *See* authentication
    CPTED, 174
    IAM. *See* identity and access management (IAM)
    media, 316
accidental threats, 71
accountability
    description, 12
    overview, 231–232
accounts
    access review, 249–250
    deprovisioning, 246–247
    registration, 246
    security process data, 274
accuracy in threat intelligence, 299
acoustical detection systems, 384
acquired software security, 416–420
acquisition phase in SDLC, 394
actors, threat, 58–59, 71–72
ad hoc networks, 199
Address Resolution Protocol (ARP), 189
administrative access in cloud-based systems, 145
administrative controls, 65
administrative investigations, 35
adversarial threats, 71
AES algorithm, 153
Agile methodology, 397–398
agreements
    contractual compliance, 25
    employment, 50–51
    SLAs, 77, 312–313, 333, 418
    third parties, 52–53
AHs (Authentication Headers), 190
AI (artificial intelligence), 336
AIKs (attestation identity keys) in TPMs, 136
ALE (annualized loss expectancy), 63
algorithms, cryptographic, 150
alignment of security functions to business requirements
    business strategy and security strategy, 17
    due care/due diligence, 20–21
    organizational processes, 18
    organizational roles and responsibilities, 18–19
    security control frameworks, 19–20
allow-listing, 327–328
amplifiers, 209
amplitude in wireless technologies, 198
analyses
    BIA. *See* business impact analysis (BIA)
    evidence life cycle, 288
    risk, 61–63
    security tests, 277–281
annualized loss expectancy (ALE), 63
annualized rate of occurrence (ARO), 63
anomaly-based intrusion analysis, 296, 332
antennas, 198
anti-malware, 334–335
appetite, risk, 60
application-layer firewalls, 210, 330
application programming interfaces (APIs) security, 421–422
applications
    description, 227
    security testing, 408–411
architectures
    asset life cycle, 106
    networks, 208–209
    security. *See* security architecture and engineering
ARO (annualized rate of occurrence), 63
ARP (Address Resolution Protocol), 189
artifacts, 289–290
artificial intelligence (AI), 336
AS (Authentication Service), 255
assemblers, 405
assembly programming language, 404
assessments
    controls, 66–67
    definition, 260–261
    disaster recovery plans, 363
    risk, 60–61
    security. *See* security assessment and testing
    software development security, 416–420
    third parties, 76
    vulnerabilities, 59, 265
assessors for security evaluations, 262–263
asset security
    classification, 86–90
    data life cycle, 99–105
    data security controls, 109–114
    handling requirements, 90–95
    management, 97–98
    objectives, 85–86
    resource provisioning, 96–99
    retention, 105–109
asset value (AV), 62
assets, 57

assisted password reset, 234
asymmetric encryption, 152–154
asynchronous attacks, 164
atomic execution, 138
attenuation in wireless technologies, 198
attestation identity keys (AIKs) in TPMs, 136
attribute-based access control (ABAC), 243
audits
    changes, 413
    cloud-based systems, 145
    conducting, 281–284
    definition, 260–261
    description, 12
    facility access, 385–386
    review and questions, 284
authentication, 229–230
    accountability, 231–232
    attribute-based access control, 243
    credential management systems, 233–234
    data communications, 221
    discretionary access control, 241
    identity proofs, 232–233
    just-in-time, 234–235
    mandatory access control, 241–242
    mechanisms, 239–244
    overview, 11–12
    pass-through, 165
    review and questions, 235–236, 243–244
    risk-based access control, 243
    role-based access control, 242
    rule-based access control, 242–243
    session management, 232
    single and multifactor, 230–231
    single sign-on, 234
Authentication Headers (AHs), 190
authentication servers, 192
Authentication Service (AS), 255
authentication systems
    Kerberos, 254–255
    Open Authorization, 253
    OpenID Connect, 253
    overview, 252–253
    remote access, 256–257
    review and questions, 257–258
    SAML, 253–254
authenticators in IEEE 802.1X, 192
authenticity in authorization, 12
authorization
    data communications, 221
    description, 12
    Open Authorization, 253
    remote access, 256–257
automating configuration management, 306
AV (asset value), 62
availability, 11, 355–356
awareness
    disaster recovery plans, 364
    methods and techniques, 80–81
    periodic content reviews, 82
    personnel security, 389
    presentation techniques, 81–82
    program effectiveness evaluation, 82
    review and questions, 82–84
    security process data, 274

B

backup power strategies, 180
backup storage
    cloud, 350
    direct-attached storage, 349
    electronic vaulting and remote journaling, 351
    network-attached storage, 350
    offline, 350
    storage area networks, 350
    strategies, 348–349
    verification data, 274
badge readers, 384
bandwidth in wireless technologies, 198
bare-metal hypervisors, 145
barriers, 380
baselines
    configuration management, 305–306
    internal governance, 42
bastion hosts, 208
BC. *See* business continuity (BC)
BCPs. *See* business continuity plans (BCPs)
behavior-based analysis
    anti-malware, 334–335
    IDS/IPS, 332
    intrusion detection, 296
Bell-LaPadula model, 126–127
BIA. *See* business impact analysis (BIA)
Biba model, 127
binding hard disk drives, 136
biometric factors in authentication, 231
bit-flipping in sandboxing, 335
black box penetration tests, 267
black-hat (unethical) hackers, 266
blacklisting, 327–328
blind penetration tests, 267
block ciphers, 152
Blowfish algorithm, 153
Bluejacking, 202
Bluesnarfing, 202
Bluetooth standard, 202
breaches
    attack simulations, 270
    data, 29–30
Brewer and Nash model, 128
bridges, 209
brouters, 209
brownouts, 180
brute force cryptanalytic attacks, 162
bus encryption, 137
business continuity (BC)
    business impact analysis, 45–47, 373–374
    introduction, 45
    overview, 372–373
    plans. *See* business continuity plans (BCPs)
    security process data, 274
business continuity plans (BCPs)
    developing, 374
    metrics, 374–375
    review and questions, 376–377
    testing, 375–376
business impact analysis (BIA)
    asset criticality, 87
    description, 46
    developing, 46
    documenting, 47

overview, 373–374
review and questions, 47–48
scope, 46–47
business strategy and security strategy, 17

C

cable modems, 216
CABs (change advisory boards), 345
California Consumer Privacy Act (CCPA), 24, 30
candidates, screening and hiring, 49
Capability Maturity Model (CMM), 399
Capability Maturity Model Integration (CMMI), 399
capacitance detectors, 385
CAs (certificate authorities), 156–157
CASBs (cloud access security brokers), 112–113
CCPA (California Consumer Privacy Act), 24, 30
CCTV (closed-circuit television) cameras, 381
CDMA (code division multiple access), 204
CDNs (content distribution networks), 147, 205–206
CEI (Computer Ethics Institute), 5
cellular networks, 204–205
Center for Internet Security (CIS) Controls, 20
certificate authorities (CAs), 156–157
certificate revocation lists (CRLs), 159
certificates, 156–157
chain of custody for evidence, 287
Challenge Handshake Authentication Protocol (CHAP), 220
change advisory boards (CABs), 345
change management
life cycle, 345–346
policies, 344–345
processes, 344–346
review and questions, 347–348
security considerations, 346–347
software development life cycle, 401
software development security, 413
channels
cellular networks, 204
communications. *See* communications channels
CHAP (Challenge Handshake Authentication Protocol), 220
chief information officers (CIOs), 19
chief information security officers (CISOs), 19
chief privacy officers (CPOs), 19
chief security officers (CSOs), 19
Chinese Wall model, 128
chosen ciphertext and chosen plaintext attacks, 163
CI/CD (continuous integration and continuous delivery), 407
CIOs (chief information officers), 19
cipher locks, 384
ciphers, 149
ciphertext, 149
ciphertext-only cryptanalytic attacks, 162
circuit-level firewalls, 210, 330
CIS (Center for Internet Security) Controls, 20
CISOs (chief information security officers), 19
civil investigations, 35–36
Clark-Wilson model, 127–128
classification
assets, 86–87
data, 87–88
information, 93–94
client-based system vulnerabilities, 140
client-to-site VPNs, 219
Clipper Chip algorithm, 150
closed-circuit television (CCTV) cameras, 381
closed-source intelligence, 300
cloud access security brokers (CASBs), 112–113
cloud-based firewalls, 330
cloud-based identity services, 238
cloud-based system vulnerabilities, 144–145
cloud sites, 354
cloud storage, 350
clustering cryptographic keys, 150
CMM (Capability Maturity Model), 399
CMMI (Capability Maturity Model Integration), 399
COBIT framework, 20
code, cryptography, 149
code, software
guidelines and standards, 420–425
repositories, 408
review and testing, 268–269
code division multiple access (CDMA), 204
code modules, 405
cohesion in programming, 405
cold sites, 353
collection
data, 102
evidence, 287–290
security data, 272–276
collisions, hash, 155
collusion, 310
columns in database tables, 141
combination locks, 384
commercial classification systems, 93
commercial-off-the-shelf (COTS) software, 416–417
commodity malware, 334
communication and network security
channels. *See* communications channels
network components, 207–214
network security application, 193–207
networking concepts, 184–192
objectives, 183–184
communication in disaster recovery plans, 361–362
communications channels
data communications, 220–222
introduction, 215
multimedia collaboration, 218–219
remote access, 219–220
review and questions, 223–224
third-party connectivity, 222–223
virtualized networks, 222
voice, 215–218
community clouds, 144
compartmented security mode, 125
compensating controls, 66
compilers, 405
compliance
contractual, 25
controls, 66–67
data security, 109–114
industry standards, 25
legal and regulatory, 24
overview, 23–24
personnel policy requirements, 53
privacy requirements, 25–26
review and questions, 26–28

compliance checks, 270
Computer Ethics Institute (CEI), 5
concentrators, 209
confidence levels in threat intelligence, 299
confidential classification, 93
confidentiality
  description, 10
  models, 126–127
configuration management
  automating, 306
  baselining, 305–306
  cloud-based systems, 145
  introduction, 304
  patches, 341–342
  provisioning, 305
  review and questions, 306–307
  software, 408
confusion in cryptographic methods, 151
constrained interfaces, 240
consultants, agreements and controls with, 52–53
containerization vulnerabilities, 146
content-dependent access authentication, 240
content distribution networks (CDNs), 147, 205–206
content reviews in security awareness, 82
context-dependent access, 240
continuous improvement in risk management, 68
continuous integration and continuous delivery (CI/CD), 407
continuous monitoring, 296
contractual compliance, 25
control planes in SDN, 196
controllers of data, 101–102
controls. *See* security controls
converged protocols, 194
copyrights, 30–31
corrective controls, 66
cost savings in third-party security services, 333
COTS (commercial-off-the-shelf) software, 416–417
countermeasure selection and implementation in risk management, 64–65
coupling in programming, 405
CPOs (chief privacy officers), 19
CPTED (Crime Prevention Through Environmental Design), 174
credentials, 11
crime prevention and disruption concerns in site planning, 167
Crime Prevention Through Environmental Design (CPTED), 174
criminal investigations, 36
critical changes, 346
criticality
  information, 131
  patch management, 340–341
CRLs (certificate revocation lists), 159
cryptanalysis, 149
cryptanalytic attacks
  brute force, 162
  chosen ciphertext and chosen plaintext, 163
  ciphertext-only, 162
  fault injection, 164
  frequency analysis, 163
  implementation, 163
  Kerberos, 165
  known plaintext, 162–163
  man-in-the-middle, 164
  overview, 161–162
  pass the hash, 165
  ransomware, 165
  review and questions, 166–167
  side-channel, 163–164
  timing, 164
cryptographic keys in TPMs, 136
cryptographic solutions
  cryptosystems, 151
  digital certificates, 156
  hybrid cryptography, 155–156
  integrity, 154–155
  introduction, 148–149
  key management practices, 158–159
  life cycle, 149–150
  methods, 151–154
  nonrepudiation and digital signatures, 158
  public key infrastructure, 156–158
  review and questions, 159–161
  vulnerabilities, 142
cryptoprocessors, 137
cryptosystems, 151
cryptovariables, 150
CSOs (chief security officers), 19
culture, risk, 60
custodians, data, 101–102
Cyber Kill Chain threat model, 72, 301
cybercrimes, 29
cycles, wave, 197

D

DAC (discretionary access control), 241
damage concerns, 168
DAST (dynamic application security testing), 410
data
  breaches, 29–30
  classification, 87–88
  communications, 220–222
  custodian responsibilities, 19
  description, 9–10
data at rest, 91, 110
data centers, 175–176
data historians in industrial control systems, 143
data in transit, 92, 109–110
data in use, 110
data life cycle
  collection, 102
  destruction, 103–104
  location, 102
  maintenance, 102–103
  overview, 99–100
  owners, 100–101
  remanence, 103
  retention, 103
  review and questions, 104–105
  roles, 100–102
data localization laws, 32
data loss prevention (DLP), 112
data owners
  cloud-based systems, 145
  responsibilities, 19
data protection requirements in control selection, 131–132
data retention in cloud-based systems, 145

data security and compliance
control scoping and tailoring, 111
control standards selection, 110–111
data at rest, 110
data in transit, 109–110
data in use, 110
data states, 109
protection methods, 111–113
review and questions, 113–114
data segmentation in cloud-based systems, 145
data sovereignty, 102
database system vulnerabilities, 141–142
DCSs (distributed control systems), 142–143
de-encapsulation in OSI model, 186
decryption, 149
dedicated security mode, 124
default-deny method of controlling access, 328
defaults
secure, 117
site and facility design, 169
defense in depth
secure design, 117
site and facility design, 169
defined maturity level in CMMI, 399
delineated responsibilities in cloud-based systems, 145
deluge water sprinkler systems, 179
demilitarized zones (DMZs)
description, 208
firewalls, 329
demotions, 51–52
deny-listing, 327–328
deprovisioning in identity and access, 245–247
DES algorithm, 153
design phase
asset life cycle, 106
software development life cycle, 394
designing evaluations, 261–264
destruction
assets, 107
data, 103–104
media, 316–317
detection
fire, 177–178
incident management, 320–321
surveillance, 381
detective and preventive controls
allow-listing and deny-listing, 327–328
anti-malware, 334–335
description, 66
firewalls, 328–330
honeypots and honeynets, 333–334
IDS/IPS, 331–332
machine learning and artificial intelligence, 336
overview, 326
review and questions, 336–338
sandboxing, 335–336
third-party provided security services, 332–333
determination, risk, 58
deterrence, surveillance for, 381
deterrent controls, 66
detonation chambers, 335
development phase in SDLC, 394
devices, 227
DevOps development, 398
DevSecOps development, 398
DHE (Diffie-Hellman) algorithm, 154
dial-up communications, 216
Diameter system, 257
Diamond Model of Intrusion Analysis, 72, 301
differential backups, 349
Diffie-Hellman (DHE) algorithm, 154
diffusion in cryptographic methods, 151
digital certificates, 156
digital forensics tools, 290–291
digital rights management (DRM), 111–112
digital signatures, 158
digital subscriber line (DSL), 216
dips, power, 180
direct-attached storage, 349
direct sequence spread spectrum (DSSS), 198
disaster recovery data in security process data, 274
disaster recovery (DR)
lessons learned, 364–365
overview, 359–360
people safety concerns, 360
plans. *See* disaster recovery plans (DRPs)
review and questions, 365–367
disaster recovery plans (DRPs)
assessment, 363
communications, 361–362
key points, 360–361
personnel, 361
response, 361
restoration, 363
testing, 367–372
training and awareness, 364
disciplinary activities, 51–52
discretionary access control (DAC), 241
disposal phase in asset life cycle, 106–107
distributed control systems (DCSs), 142–143
distributed system vulnerabilities, 141
distribution facilities, 175
DLP (data loss prevention), 112
DMZs (demilitarized zones)
description, 208
firewalls, 329
documentation
business impact analysis, 47
investigations, 292–293
software configuration management, 408
dogs, 381–382
door delay locks, 384
double blind penetration tests, 267
DR. *See* disaster recovery (DR)
drives
binding, 136
self-encrypting, 137
DRM (digital rights management), 111–112
DRPs. *See* disaster recovery plans (DRPs)
dry pipe water sprinkler systems, 179
DSL (digital subscriber line), 216
DSSS (direct sequence spread spectrum), 198
due care, 20–21
due diligence, 20–21
duress codes for locks, 384
duress systems for personnel safety, 390–391
duties, separation of
information security, 118
overview, 14
purpose, 309–310
site and facility design, 170
dynamic application security testing (DAST), 410

E

EAP (Extensible Authentication Protocol), 191–192, 220
EAR (Export Administration Regulation), 31
ECC (elliptic curve cryptography), 154
Economic Espionage Act, 30
edge computing system vulnerabilities, 146–147
EDLP (Endpoint DLP), 112
education for security awareness, 80–84
EF (exposure factor) in risk, 62–63
effectiveness
    controls, 66–67
    software development security, 412–415
egress monitoring, 297–298
EKs (endorsement keys) in TPMs, 136
El Gamal algorithm, 154
electromagnetic (EM) spectrum, 197
electromechanical intrusion detection systems, 385
electronic locks, 384
electronic vaulting, 351
elliptic curve cryptography (ECC), 154
EM (electromagnetic) spectrum, 197
embedded system vulnerabilities, 143–144
emergency changes, 346
emergency management, 389–390
emergency procedures, 360
employment agreements and policies, 50
Encapsulating Security Payload (ESP), 190
encapsulation
    network security, 196–197
    OSI model, 186
encryption. *See also* cryptographic solutions
    bus, 137
    data communications, 221
    description, 149
end-of-life phase in asset life cycle, 106–107
end-of-support phase in asset life cycle, 106–107
endorsement keys (EKs) in TPMs, 136
Endpoint DLP (EDLP), 112
endpoint security for network components, 213
engineering security. *See* security architecture and engineering
enrollment for user accounts, 246
entities, 9–10
entry control points, 379
environmental issues
    site and facility, 177
    threats, 71
ESP (Encapsulating Security Payload), 190
establishing identity, 232–233
ethical disclosure in security test results, 279
ethical (white-hat) hackers, 266
ethics. *See* professional ethics
evacuation plans, 179–180
evaluations
    assessors, 262–263
    goals and strategies, 261–262
    review and questions, 263–264
event logs, 298
evidence
    collection and handling, 287–290
    life cycle, 287–288
    storage, 176
evolutionary prototypes, 396
exception handling in security test results, 279
explicit denies, 328
explicit permissions, 241
Export Administration Regulation (EAR), 31
exposure factor (EF) in risk, 62–63
Extensible Authentication Protocol (EAP), 191–192, 220
external assessors for security evaluations, 262–263
external governance, 16
external security auditors, 282–283
extranets, 208
Extreme Programming (XP), 398

F

facilities
    access audits, 385–386
    description, 227
    design. *See* site and facility design
factors, authentication, 230–231
fail securely
    secure design, 117–118
    site and facility design, 169–170
fault injection cryptanalytic attacks, 164
fault tolerance, 356
faults, power, 180
FCIP (Fibre Channel over IP), 194
FCoE (Fibre Channel over Ethernet), 194
FDMA (frequency division multiple access), 204
Federal Information Security Management Act (FISMA), 24
Federated Identity Management (FIM), 233, 237–239
fencing, 379–380
FHSS (frequency hopping spread spectrum), 198
Fibre Channel over Ethernet (FCoE), 194
Fibre Channel over IP (FCIP), 194
fields in database tables, 141
fifth generation programming languages, 404
FIM (Federated Identity Management), 233, 237–239
fire prevention, 177–180
Firewall as a Service (FWaaS), 330
firewalls
    cloud-based, 330
    network- and host-based, 329
    overview, 328–329
    types, 210–211, 329–330
    web application, 330
firmware in system security, 135
first generation programming languages, 404
FISMA (Federal Information Security Management Act), 24
flat-file databases, 142
footprints of satellites, 203
foreign keys in database systems, 141–142
forensic investigations
    description, 287
    digital tools, 290–291
forwarding planes in SDN, 196
fourth generation programming languages, 404
frameworks, risk, 64
frequencies in wireless technologies, 197–198
frequency analysis cryptanalytic attacks, 163
frequency bands in wireless technologies, 198
frequency division multiple access (FDMA), 204
frequency hopping spread spectrum (FHSS), 198
full backups, 349
full interruption tests
    business continuity plans, 375
    disaster recovery plans, 370

full-knowledge penetration tests, 267
functional requirements in control selection, 131
functions of controls, 65–66
fuzzing, 411
FWaaS (Firewall as a Service), 330

G

gates, 380
gateways, 209
General Data Protection Regulation (GDPR), 24, 101–102
generators, power, 180
geosynchronous satellites, 203
goals
    organizational, 17
    security evaluations, 261–262
governance requirements
    control selection, 132
    ethics, 6
Gramm-Leach-Bliley Act, 24, 30
gray box penetration tests, 267
gray-hat hackers, 266
guards, 381–382
guidelines for internal governance, 41–42

H

H.323 standard, 217
HA (high availability), 355–356
handling assets
    classification systems, 93–94
    overview, 90–91
    review and questions, 94–95
    storage, 91–92
    transfer, 93
    transmission, 92
    transportation, 92
handling evidence, 287–290
hard disk drives, binding, 136
hardware
    network components, 208–209
    risks, 75
    system security, 135
hardware security modules (HSMs), 136
harm to people, preventing, 360
hashes in cryptography, 155
Health Information Technology for Economic and Clinical Health (HITECH) Act, 24, 30
Health Insurance Portability and Accountability Act (HIPAA), 24
heating, ventilation, and air conditioning (HVAC) services, 177
heuristic analysis
    IDS/IPS, 332
    intrusion detection, 296
heuristic anti-malware, 334–335
HIDS/HIPS (host-based intrusion detection/prevention system), 331
hierarchical databases, 142
high availability (HA), 355–356
high-level programming languages, 404
high-performance computing (HPC) system vulnerabilities, 146
HIPAA (Health Insurance Portability and Accountability Act), 24
hiring candidates, 49
HITECH (Health Information Technology for Economic and Clinical Health) Act, 24, 30
HMIs (human–machine interfaces), 143
honeypots and honeynets, 333–334
horizontal enactments, 33
host-based firewalls, 329
host-based intrusion detection/prevention system (HIDS/HIPS), 331
hot sites, 353
HPC (high-performance computing) system vulnerabilities, 146
HSMs (hardware security modules), 136
HTTP (Hypertext Transfer Protocol), 189–190
hubs, 209
human–machine interfaces (HMIs), 143
HVAC (heating, ventilation, and air conditioning) services, 177
hybrid clouds, 144
hybrid cryptography, 155–156
hybrid identity services, 238
Hypertext Transfer Protocol (HTTP), 189–190
hypervisors, 145

I

IaaS (Infrastructure as a Service), 144
IAM. *See* identity and access management (IAM)
ICMP (Internet Control Message Protocol), 188
ICSs (industrial control systems), 142–143
IDE (integrated development environment), 406
identification
    data communications, 220
    description, 11
    threats and vulnerabilities, 59–60
identity and access management (IAM)
    authentication systems, 252–258
    authorization mechanisms, 239–244
    federated identity, 237–239
    identification and authentication, 229–236
    logical and physical access control, 226–229
    objectives, 225–226
    provisioning life cycle, 245–252
Identity management (IdM), 230
identity providers in SAML, 254
IDFs (intermediate distribution facilities), 175
IdM (Identity management), 230
IDS/IPSs (intrusion detection and prevention systems)
    categories, 331–332
    models, 296
    physical IDSS, 384–385
IEEE (Institute of Electrical and Electronics Engineers)
    802.1X, 192, 220
    Code of Ethics, 6
IGMP (Internet Group Management Protocol), 188
IKE (Internet Key Exchange), 190
impact, risk, 58
implementation
    asset life cycle, 106
    change, 345
implementation cryptanalytic attacks, 163
implicit denies, 328
implicit permissions, 241
import/export controls, 31
important changes, 346
incident concerns in site and facility design, 168

incident management
detection, 320–321
lessons learned, 324
life cycle, 319–320
mitigation, 322
overview, 318–319
preparation, 320
recovery, 323
remediation, 323
reporting, 322–323
response, 321–322
review and questions, 324–326
incomplete maturity level in CMMI, 399
incremental backups, 349
Incremental methodology, 396
indexes in database systems, 142
indicators of compromise (IoCs)
incident management, 320–321
threats, 71
industrial control systems (ICSs), 142–143
industry standards
compliance, 25
investigations, 37
information, description, 9–10, 227
information flow models, 127
information systems (IS)
capabilities, 135–138
review and questions, 138–139
Infrastructure as a Service (IaaS), 144
infrastructure mode in Wi-Fi, 199
initial maturity level in CMMI, 399
initial response phase in evidence life cycle, 288
initiators in iSCSI, 195
inrush current, power, 180
Institute of Electrical and Electronics Engineers (IEEE)
802.1X, 192, 220
Code of Ethics, 6
intangible assets, 57, 97
integrated development environment (IDE), 406
integrated product teams (IPTs), 401
Integrated Services Digital Network (ISDN), 216
integration tests in application security, 409
integrity
cryptography, 154–155
data communications, 221
overview, 11
integrity models, 127–128
intellectual property (IP) requirements, 30–31
interface requirements in control selection, 132
interface tests
application security, 409
security issues, 269–270
interference in wireless technologies, 198
intermediate CAs, 157
intermediate distribution facilities (IDFs), 175
internal assessors for security evaluations, 262–263
internal governance
baselines, 42
description, 16–17
guidelines, 41–42
overview, 39–40
policies, 40
procedures, 40–41
review and questions, 42–44
standards, 41
internal security auditors, 282
internal security controls
locks, 383–384
overview, 382–383
personnel entry requirements, 383
International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 27002, 20
International Traffic in Arms Regulations (ITAR), 31
Internet Control Message Protocol (ICMP), 188
Internet Group Management Protocol (IGMP), 188
Internet Key Exchange (IKE), 190
Internet of Things (IoT) vulnerabilities, 143
Internet Protocol (IP), 187–188
Internet Protocol (IP) telephony, 217
Internet Security Association and Key Management Protocol (ISAKMP), 190
Internet Small Computer Systems Interface (iSCSI), 194–195
interpreters, 405
intranets, 208
intrusion detection and prevention systems (IDS/IPSs)
categories, 331–332
models, 296
physical IDSS, 384–385
inventory of assets, 96–97
investigations
administrative, 35
civil, 35–36
criminal, 36
digital forensics tools, 290–291
evidence collection and handling, 287–290
forensic, 287
industry standards, 37
overview, 286
regulatory, 36–37
reporting and documentation, 292–293
review and questions, 37–39, 293–294
techniques, 291–292
invocation rule in Biba model, 127
IoCs (indicators of compromise)
incident management, 320–321
threats, 71
IoT (Internet of Things) vulnerabilities, 143
IP (intellectual property) requirements, 30–31
IP (Internet Protocol), 187–188
IP (Internet Protocol) telephony, 217
IP Security (IPSec), 190–191
IPTs (integrated product teams), 401
IS (information systems)
capabilities, 135–138
review and questions, 138–139
ISAKMP (Internet Security Association and Key Management Protocol), 190
(ISC)² Code of Ethics
canons, 3–4
preamble, 3
iSCSI (Internet Small Computer Systems Interface), 194–195
ISDN (Integrated Services Digital Network), 216
ISO (International Organization for Standardization)/International Electrotechnical Commission (IEC) 27002, 20
ITAR (International Traffic in Arms Regulations), 31
iterative methodology, 396

J

job rotation, 311–312
journaling, remote, 351
just-in-time access control, 234–235

K

Kanban methodology, 398
KCIs (key control indicators), 275
KDCs (Key Distribution Centers), 254–255
keep it simple principle
    secure design, 119
    site and facility design, 170
Kerberos
    characteristics, 220
    components, 254–255
    cryptanalytic attacks, 165
    description, 192
Kerckhoffs, Auguste, 150
Kerckhoffs' principle, 150
key control indicators (KCIs), 275
Key Distribution Centers (KDCs), 254–255
key goal indicators (KGIs), 275
key management practices, 158–159
key override for locks, 384
key performance indicators (KPIs), 275
key risk indicators (KRIs), 275
keys
    asymmetric encryption, 152–153
    cryptographic, 150
    database systems, 141–142
    hybrid cryptography, 156
    symmetric encryption, 151–152
    Zigbee, 203
KGIs (key goal indicators), 275
knowledge-based intrusion detection, 332
known plaintext cryptanalytic attacks, 162–163
KPIs (key performance indicators), 275
KRIs (key risk indicators), 275

L

Layer 2 Tunneling Protocol (L2TP), 221
layered security in site and facility design, 169
least privilege principle
    description, 13
    secure design, 116–117
    security operations, 309
    site and facility design, 169
legal and regulatory requirements
    cybercrimes, 29
    data breaches, 29–30
    import/export controls, 31
    licensing and intellectual property, 30–31
    privacy issues, 32–33
    review and questions, 33–34
    transborder data flow, 32
legal compliance, 24
legal holds, 289
legal liability
    cloud-based systems, 145
    third-party provided security services, 333
LEO (Low Earth orbit) satellites, 203
lessons learned
    disaster recovery, 364–365
    incident management, 324
Li-Fi technology, 203–204
libraries, software, 405–406
licensing requirements, 30–31
life cycle
    assets, 106
    change management, 345–346
    cryptographic, 149–150
    data, 99–105
    evidence, 287–288
    identity and access provisioning, 245–252
    incident management, 319–324
    software development, 394–395
light in intrusion detection systems, 384
lighting, 380–381
likelihood, risk, 58
link keys in Zigbee, 203
lives, saving, 360
location of data, 102
locks, 383–384
log management, 298
log reviews, 267–268
logging and monitoring
    changes, 413
    continuous monitoring, 296
    egress monitoring, 297–298
    intrusion detection and prevention, 296
    log management, 298
    overview, 295
    review and questions, 302–304
    SIEM systems, 297
    threat hunting, 300
    threat intelligence, 298–300
    threat modeling, 300–301
    user and entity behavior analytics, 301–302
logical controls
    description, 65
    IAM, 226–229
loose coupling in programming, 405
Low Earth orbit (LEO) satellites, 203
lumens, 380

M

m-of-n control, 310
MAC (mandatory access control), 122, 241–242
machine learning (ML), 336
machine programming languages, 404
main distribution facilities (MDFs), 175
maintenance
    asset life cycle, 106
    CPTED, 174
    data, 102–103
    software development life cycle, 395
malware, 334–335
man-in-the-middle (MITM) attacks, 164
managed maturity level in CMMI, 399
managed security services (MSSs), 332
managed service accounts, 249
managed services, 418
management review and approval for security process data, 275
managerial controls, 65
mandatory access control (MAC), 122, 241–242
mandatory vacations, 312
master keys for locks, 384
master keys in Zigbee, 203

maturity models, 398–399
maximum tolerable downtime (MTD), 374
MDFs (main distribution facilities), 175
mean time between failures (MTBF), 375
mean time to failure (MTTF), 375
mean time to repair (MTTR), 375
measuring risk, 67
mechanical locks, 383–384
media management, 315–317
media storage facilities, 176
memory protection in trusted execution environments, 137
message digests in cryptography, 155
metrics in business continuity plans, 374–375
MFA (multifactor authentication), 230–231
micro-segmentation, 195
microservices vulnerabilities, 146
Microsoft MS-CHAPv2, 220
minimum security requirements in supply chain risk management, 77
missions, organizational, 17
misuse case testing, 269
mitigation
    incident management, 322
    software development security, 413–414
MITM (man-in-the-middle) attacks, 164
MITRE ATT&CK Framework, 72, 301
ML (machine learning), 336
mobile code, 406
mobile sites, 355
modulation in wireless technologies, 198
monitoring
    and logging. *See* logging and monitoring
    risk, 67
    third parties, 76
motion in intrusion detection systems, 385
MSSs (managed security services), 332
MTBF (mean time between failures), 375
MTD (maximum tolerable downtime), 374
MTTF (mean time to failure), 375
MTTR (mean time to repair), 375
multicasting in IGMP, 188
multifactor authentication (MFA), 230–231
multilayer protocols, 186, 193–197
multilevel security mode, 125–126
multimedia collaboration, 218–219
multiperson control, 309–310
multiple processing sites, 352
multistate systems, 124

N

NAC (network access control) devices, 211
NAS (network-attached storage), 350
National Institute of Standards and Technology (NIST)
    incident response life cycle, 319
    Risk Management Framework, 61
    Special Publication 800-18, 100
    Special Publication 800-37, 100
    Special Publication 800-53, 20, 26
natural access control in CPTED, 174
natural programming languages, 404
NDLP (Network DLP), 112
need-to-know principle
    authentication, 240
    description, 13
    security models, 123
    security operations, 308
network access control (NAC) devices, 211
network-attached storage (NAS), 350
network-based firewalls, 329
network-based IDS/IPS (NIDS/NIPS), 331
network components
    controls, 211–212
    endpoint security, 213
    firewalls, 210–211
    hardware, 208–209
    network access control devices, 211
    review and questions, 214
    transmission media, 212–213
Network DLP (NDLP), 112
network keys in Zigbee, 203
network security application
    content distribution networks, 205–206
    introduction, 193
    multilayer protocols, 193–197
    review and questions, 206–207
    wireless technologies, 197–205
networking concepts
    ARP, 189
    EAP, 191–192
    ICMP, 188
    IEEE 802.1X, 192
    IGMP, 188
    Internet Protocol, 187–188
    introduction, 184
    IPSec, 190–191
    Kerberos, 192
    OSI model, 185–187
    secure protocols, 189
    Secure Shell, 191
    SSL, TLS, and HTTPS, 189–190
    TCP/IP stack, 186–188
next-generation firewalls (NGFWs), 210, 330
NIDS/NIPS (network-based IDS/IPS), 331
NIST. *See* National Institute of Standards and Technology (NIST)
nonces, 151
noninterference model, 128
nonrepudiation
    cryptography, 158
    data communications, 221
    description, 12–13
nontechnical vulnerabilities, 340
NoSQL database systems, 142

O

O&M (operation and maintenance) in software development life cycle, 400
OAuth (Open Authorization), 253
objects
    entities, 10
    IAM, 226
    security models, 122
OCSP (Online Certificate Status Protocol), 159
OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation) threat model, 72, 301
OFDM (orthogonal frequency division multiplexing), 198, 204
OFDMA (orthogonal frequency division multiple access), 204

offline storage, 350
OIDC (OpenID Connect), 253
on-path attacks, 164
on-premise identity management, 237–238
onboarding, 51
one-way cryptographic functions, 151
Online Certificate Status Protocol (OCSP), 159
Open Authorization (OAuth), 253
open-source intelligence (OSINT), 299–300
open-source software, 417
Open Systems Interconnection (OSI) model, 185–187
open trust model in Zigbee, 203
Open Web Application Security Project (OWASP), 400
OpenID Connect (OIDC), 253
OpenIOC threat model, 301
operation and maintenance (O&M) in software development life cycle, 400
operational controls, 65
operational goals, organizational, 17
operational prototypes, 396
Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) threat model, 72, 301
operations phase in software development life cycle, 395
optimizing maturity level in CMMI, 399
organizational code of ethics, 4
  review and questions, 7–8
  sources, 5–6
  workplace ethics statements and policies, 4–6
organizational processes, 18
organizational roles and responsibilities, 18–19
orthogonal frequency division multiple access (OFDMA), 204
orthogonal frequency division multiplexing (OFDM), 198, 204
OSI (Open Systems Interconnection) model, 185–187
OSINT (open-source intelligence), 299–300
OWASP (Open Web Application Security Project), 400
owners
  assets, 96
  cloud-based systems, 145
  data life cycle, 100–101
  responsibilities, 19

P

PaaS (Platform as a Service), 144
packet-filtering firewalls, 210, 329
padlocks, 383
pair programming, 398
pairing in Bluetooth, 202
PAP (Password Authentication Protocol), 220
parallel testing
  business continuity plans, 375
  disaster recovery plans, 370
partial-knowledge penetration tests, 267
pass the hash cryptanalytic attacks, 165
passive infrared systems, 384
Password Authentication Protocol (PAP), 220
password management and synchronization, 234
PASTA (Process for Attack Simulation and Threat Analysis), 72, 301
patch management, 339–340
  cloud-based systems, 145
  criticality, 340–341
  review and questions, 342–344
  schedules, 341
  testing, 341–342
patents, 30–31
pathping utility, 188
pattern-based anti-malware, 334–335
pattern-based intrusion detection, 296, 332
Payment Card Industry (PCI) Data Security Standards (DSS), 20
PBX (Private Branch Exchange) systems, 215–216
PCRs (platform configuration registers), 136
penetration testing
  application security, 410
  security control, 265–267
people safety concerns in disaster recovery, 360
performance requirements in control selection, 131
Perimeter Intrusion Detection and Assessment Systems (PIDASs), 379–380
perimeter security controls, 378
  barriers, 380
  entry control points, 379
  fencing, 379–380
  guards and dogs, 381–382
  lighting, 380–381
  surveillance, 381
  zones, 378–379
periodic content reviews for security awareness, 82
permissions, 240–241
persistent memory in TPM, 136
personnel
  communications, 361
  entry requirements, 383
  privacy policy requirements, 53–54
personnel safety
  duress, 390–391
  emergency management, 389–390
  review and questions, 391–392
  security training and awareness, 389
  travel, 388–389
personnel security, 48
  candidate screening and hiring, 49
  compliance policy requirements, 53
  employment agreements and policies, 50
  onboarding, 51
  practices, 49
  review and questions, 54–56
  terminations, 52
  third parties, 52–53
  transfers, promotions, and disciplinary activities, 51–52
photometric sensors, 384
physical security
  description, 65
  facility access audits, 385–386
  IAM, 226–229
  internal controls, 382–387
  intrusion detection systems, 384–385
  overview, 377–378
  perimeter security controls, 378–382
  review and questions, 386–387
physical Wi-Fi standards, 199–201
PIDASs (Perimeter Intrusion Detection and Assessment Systems), 379–380
ping utility, 188
PKI (public key infrastructure), 156–158
plain old telephone service (POTS), 215
plaintext, 149

Plan of Action and Milestones (POA&M), 67
Platform as a Service (PaaS), 144
platform configuration registers (PCRs), 136
PLCs (programmable logic controllers), 143
POA&M (Plan of Action and Milestones), 67
Point-to-Point Protocol (PPP), 221
Point-to-Point Tunneling Protocol (PPTP), 221
poisoning, ARP, 189
policies
    change management, 344–345
    compliance, 53
    employment, 50
    internal governance, 40
    privacy, 26, 53–54
    workplace ethics, 4–5
portable code, 406
post-change activities, 345
POTS (plain old telephone service), 215
power for site and facility, 180
PPP (Point-to-Point Protocol), 221
PPTP (Point-to-Point Tunneling Protocol), 221
preaction water sprinkler systems, 179
preparation in incident management, 320
presentation
    evidence life cycle, 288
    security awareness, 81–82
preventive controls, 66
primary keys in database systems, 141
principals
    Kerberos, 254
    SAML, 254
principle of least privilege
    description, 13
    secure design, 116–117
    security operations, 309
    site and facility design, 169
privacy
    compliance, 25–26
    legal and regulatory requirements, 32–33
    personnel requirements, 53–54
    policy requirements, 53–54
privacy by design
    secure design, 119
    site and facility design, 171
Private Branch Exchange (PBX) systems, 215–216
private clouds, 144
private keys in asymmetric encryption, 152–153
private organizations in classification systems, 93
privilege creep, 51
privilege escalation, 248–249
privileged account management, 310–311
privileges, 240–241
procedures in internal governance, 40–41
Process for Attack Simulation and Threat Analysis (PASTA), 72, 301
processing modes in security models, 124–126
processor security extensions, 138
processors, data, 101–102
professional ethics
    introduction, 2
    (ISC)² Code of Ethics, 3–4
    organizational, 4–6
program effectiveness evaluation for security awareness, 82
programmable logic controllers (PLCs), 143
programming languages, 404–405
promotions, 51–52
proofing identity, 232–233
proprietary classification, 93
proprietary intelligence, 300
protection methods for data, 111–113
prototyping methodology, 396
provisioning
    configuration management, 305
    resources, 96–99
provisioning in identity and access
    account access review, 249–250
    deprovisioning, 245–247
    initial, 245–246
    managed service accounts, 249
    overview, 245
    privilege escalation, 248–249
    review and questions, 251–252
    role definitions, 247–248
proxies, 209
proximity badges, 384
proximity detectors, 385
proxy firewalls, 210, 330
PSTNs (public switched telephone networks), 215
public classification, 93
public clouds, 144
public key infrastructure (PKI), 156–158
public keys in asymmetric encryption, 152–153
public switched telephone networks (PSTNs), 215

Q

QKD (quantum key distribution), 154
qualitative risk analysis, 62
quality of service (QoS), 356
quantitative risk analysis, 62
quantitatively managed maturity level in CMMI, 399
quantum cryptography, 154
quantum key distribution (QKD), 154

R

RAD (Rapid Application Development), 397
RADIUS (Remote Authentication Dial-In User Service), 221, 256
ransomware, 165
Rapid Application Development (RAD), 397
rapid prototyping, 396
RAs (registration authorities) for certificates, 157
ratings for threat intelligence, 299
RBAC (role-based access control), 242, 248
RC4 algorithm, 153
read-through tests
    business continuity plans, 375
    disaster recovery plans, 368–369
Real-time Transport Protocol (RTP), 217
realms, Kerberos, 254
reciprocal sites, 354
records in database tables, 141
recovery
    controls, 66
    incident management, 323
recovery point objective (RPO), 375
recovery sites
    cloud, 354
    cold, 353
    hot, 353

mobile, 355
multiple processing, 352
reciprocal, 354
strategies, 351–352
warm, 353
recovery strategies
backup storage, 348–351
fault tolerance, 356
high availability, 355–356
quality of service, 356
recovery sites, 351–355
resiliency, 355
review and questions, 357–359
recovery time objective (RTO), 374
reflection in wireless technologies, 198
registration
identity, 232–233
user accounts, 246
registration authorities (RAs) for certificates, 157
regulatory compliance, 24
regulatory investigations, 36–37
relational databases, 141–142
relevance in threat intelligence, 299
relying parties in OpenID Connect, 253
remanence, data, 103
remediation
incident management, 323
security test results, 278–279
remote access
authentication and authorization, 256–257
overview, 219–220
Remote Authentication Dial-In User Service (RADIUS), 221, 256
remote journaling, 351
repeaters, 209
replacement phase
asset life cycle, 106–107
software development life cycle, 395
reporting
incident management, 322–323
investigations, 292–293
risk, 67–68
security tests, 277–278
reputation-based anti-malware, 334–335
requests for change, 345
residual risk, 64
resiliency, 355
resource protection
media management, 315–317
overview, 314–315
review and questions, 317–318
resource provisioning, 96–99
response
disaster recovery plans, 361
incident management, 321–322
risk, 63–64
responsibilities
alignment of security functions to business requirements, 18–19
cloud-based systems, 145
restoration in disaster recovery plans, 363
restricted area security, 176
retention
assets, 105–109
cloud-based systems, 145
data, 103
retirement phase
asset life cycle, 106
software development life cycle, 395
rights, 240–241
risk analysis in software development security, 413–414
risk-based access control, 243
risk management
continuous improvement, 68
controls, 65–67
countermeasure selection and implementation, 64–65
monitoring, 66–67
reporting, 67–68
review and questions, 68–70
risk analysis, 61–63
risk assessment, 60–61
risk elements, 57–58
risk frameworks, 64
risk response, 63–64
threats and vulnerabilities identification, 59–60
Risk Management Framework (RMF), 61
risk registers, 67
risk response in control selection, 133
risk sharing in third-party provided security services, 333
Rivest, Shamir, Adleman (RSA) algorithm, 154
RMF (Risk Management Framework), 61
role-based access control (RBAC), 242, 248
roles
alignment of security functions to business requirements, 18–19
data, 100–102
IAM, 247–248
root CAs, 157
routers, 209
routine changes, 346
rows in database tables, 141
RPO (recovery point objective), 375
RSA (Rivest, Shamir, Adleman) algorithm, 154
RTO (recovery time objective), 374
RTP (Real-time Transport Protocol), 217
rule-based access control, 242–243
rule-based intrusion detection, 296, 332
rules in Bell-LaPadula model, 126–127
runtime environments, 406–407

S

SaaS (Software as a Service), 144
safety, personnel, 388–392
sags, power, 180
SAML (Security Assertion Markup Language), 253–254
SAMM (Software Assurance Maturity Model), 400
sandboxing, 335–336
sanitizing media, 316–317
SANs (storage area networks), 350
Sarbanes-Oxley (SOX) Act, 24
SAST (static application security testing), 410
satellites in wireless technologies, 203
saving lives, 360
SCADA (supervisory control and data acquisition) systems, 142–143
schedules for patch updates, 341
SCM (software configuration management), 408
scope in business impact analysis, 46–47
scoping and tailoring controls, 111
screened subnets, 329

screening candidates, 49
screening routers, 210, 329
SCRM. *See* supply chain risk management (SCRM)
Scrum methodology, 397
SD-WAN (software-defined wide area networking), 196, 222
SDLC. *See* software development life cycle (SDLC)
SDN (software-defined networking), 195–196, 222
sealing process in TPM, 136
second generation programming languages, 404
secret keys
    asymmetric encryption, 152–154
    symmetric encryption, 151–152
secure defaults, 117
secure processing, 137
secure protocols, 189
Secure Real-time Transport Protocol (SRTP), 217
Secure Shell (SSH), 191, 221
Secure Sockets Layer/ Transport Layer Security (SSL/TLS), 221
security and risk management
    business continuity, 45–48
    compliance, 23–28
    internal governance, 39–44
    investigations, 35–39
    legal and regulatory requirements, 29–34
    objectives, 1–2
    personnel security policies and procedures, 48–56
    professional ethics, 2–8
    risk management, 57–70
    security awareness, 80–84
    security concepts, 9–11
    security governance principles, 16–22
    supply chain risk management, 74–79
    tenets of information security, 11–16
    threat modeling, 70–74
security architecture and engineering
    control selection, 130–134
    cryptanalytic attacks, 161–167
    cryptographic solutions, 148–161
    defense in depth, 117
    fail securely, 117–118
    information systems, 135–139
    least privilege principle, 116–117
    objectives, 115–116
    principles, 115–122
    privacy by design, 119
    secure defaults, 117
    security models, 122–130
    separation of duties, 118
    shared responsibility, 120
    simplicity, 119
    site and facility controls, 173–181
    site and facility design, 167–172
    threat modeling, 116
    trust but verify principle, 119–120
    vulnerabilities, 139–148
    zero trust, 119
Security as a Service, 332
Security Assertion Markup Language (SAML), 253–254
security assessment and testing
    analyses, 277–281
    audits, 281–284
    definitions, 260–261
    evaluations, 261–264
    external auditors, 282–283
    internal auditors, 282
    objectives, 259–260
    security control testing, 264–272
    security process data collection, 272–276
    third-party auditors, 283–284
security auditor responsibilities, 19
security awareness
    methods and techniques, 80–81
    periodic content reviews, 82
    presentation techniques, 81–82
    program effectiveness evaluation, 82
    review and questions, 82–84
security clearances, 240
security concepts
    availability, 11
    confidentiality, 10
    data, information, systems, and entities, 9–10
    integrity, 11
    introduction, 9
    tenets of information security, 13–14
security control frameworks, 19–20
security control selection
    data protection requirements, 131–132
    governance requirements, 132
    interface requirements, 132
    overview, 130–131
    performance and functional requirements, 131
    review and questions, 133–134
    risk response requirements, 133
    standards selection, 110–111
security control testing
    breach attack simulations, 270
    code review and testing, 268–269
    compliance checks, 270
    interface testing, 269–270
    log reviews, 267–268
    misuse cases, 269
    overview, 264
    penetration testing, 265–267
    review and questions, 271–272
    synthetic transactions, 268
    test coverage analysis, 269
    vulnerability assessments, 265
security controls
    assessments, 66–67
    data security, 109–114
    functions, 65–66
    network, 211–212
    site and facility, 173–181
    software development security, 403–412
    third parties, 52–53
    types, 65
security governance
    alignment of security functions to business requirements, 17–21
    external, 16
    internal, 16–17
    introduction, 16
    review and questions, 21–22
security information and event management (SIEM) systems, 297
security models, 122
    confidentiality models, 126–127
    integrity models, 127–128

review and questions, 128–130
system states and processing modes, 124–126
terms and concepts, 123–124
security operations
business continuity, 372–377
change management, 344–348
concepts, 308–314
configuration management, 304–307
detective and preventive controls, 326–338
disaster recovery, 359–367
incident management, 318–326
investigations, 286–294
job rotation, 311–312
logging and monitoring, 295–304
mandatory vacations, 312
need-to-know principle, 308
objectives, 285–286
patch and vulnerability management, 339–344
physical security, 377–387
principle of least privilege, 309
privileged account management, 310–311
recovery strategies, 348–359
resource protection, 314–318
review and questions, 313–314
separation of duties, 309–310
service level agreements, 312–313
security orchestration, automation, and response (SOAR), 407
security process data, collecting, 272–276
security services by third parties, 332–333
security strategy and business strategy, 17
security tests for application security, 409
security through obscurity, 150
self-encrypting drives (SEDs), 137
self-service password reset, 234
sensitive classification, 93
sensitivity levels in security models, 122
sensitivity of information, 131
separation of duties (SoD)
information security, 118
overview, 14
purpose, 309–310
site and facility design, 170
server-based system vulnerabilities, 140–141
server rooms, 175–176
serverless function vulnerabilities, 146
service level agreements (SLAs)
security operations, 312–313
supply chain risk management, 77
third-party provided security services, 333
third-party software, 418
service providers in SAML, 254
service risks, 75–76
Service Set Identifiers (SSIDs), 199
Session Initiation Protocol (SIP), 217
session keys
hybrid cryptography, 156
symmetric encryption, 151–152
session management, 232
shared responsibility
secure design, 120
site and facility design, 171–172
side-channel cryptanalytic attacks, 163–164
SIEM (security information and event management) systems, 297
signaling in wireless technologies, 197–198
signature-based anti-malware, 334–335
signature-based intrusion detection, 296, 332
signatures in cryptography, 158
simple integrity rule in Biba model, 127
simple security rule in Bell-LaPadula model, 126
simplicity
design principle, 119
site and facility design, 170
simulations
breach attacks, 270
business continuity plans, 375
disaster recovery plans, 369–370
sine waves in wireless technologies, 197
single-factor authentication, 230–231
single loss expectancy (SLE), 62–63
single sign-on (SSO), 234
single-state systems, 124
SIP (Session Initiation Protocol), 217
site and facility controls
areas of concern, 174–176
CPTED, 174
fire, 177–180
overview, 173–174
power, 180
review and questions, 181–182
utilities, 177
site and facility design
defaults, 169
defense in depth, 169
fail securely, 169–170
least privilege, 169
privacy by design, 171
review and questions, 172–173
separation of duties, 170
shared responsibility, 171–172
simplicity, 170
site planning, 167–168
threat modeling, 168–169
trust but verify principle, 171
zero trust principle, 171
site recovery. *See* recovery sites
site-to-site VPNs, 219
Skipjack algorithm, 150
SLAs. *See* service level agreements (SLAs)
SLE (single loss expectancy), 62–63
SOAR (security orchestration, automation, and response), 407
SoD. *See* separation of duties (SoD)
Software as a Service (SaaS), 144
Software Assurance Maturity Model (SAMM), 400
software configuration management (SCM), 408
software-defined networking (SDN), 195–196, 222
software-defined security, 424
software-defined wide area networking (SD-WAN), 196, 222
software development life cycle (SDLC)
change management, 401
integrated product teams, 401
life cycle, 394–395
maturity models, 398–400
methodologies, 395–398
operation and maintenance, 400
review and questions, 401–403

software development security
application security testing, 408–411
assets, 416–420
code repositories, 408
coding guidelines and standards, 420–425
continuous integration and continuous delivery, 407
controls, 403–412
DAST, 410
effectiveness, 412–415
integrated development environment, 406
libraries, 405–406
objectives, 393–394
programming languages, 404–405
runtime, 406–407
SAST, 410
SOAR, 407
software configuration management, 408
software development life cycle, 394–403
tool sets, 406
software risks, 75
software version control, 401
sound in intrusion detection systems, 384
source-code level weaknesses and vulnerabilities, 420–421
sources for security process data, 273
SOX (Sarbanes-Oxley) Act, 24
spikes, power, 180
Spiral developmental methodology, 396
Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege (STRIDE) threat model, 72, 301
spread-spectrum wireless technology, 198
sprints, 397
SRKs (storage root keys) in TPM, 136
SRTP (Secure Real-time Transport Protocol), 217
SSH (Secure Shell), 191, 221
SSIDs (Service Set Identifiers), 199
SSL/TLS (Secure Sockets Layer/ Transport Layer Security), 221
SSO (single sign-on), 234
standards
control, 110–111
industry, 25, 37
internal governance, 41
NIST. *See* National Institute of Standards and Technology (NIST)
software code, 420–425
Wi-Fi, 199–201
star integrity rule in Biba model, 127
star property rule in Bell-LaPadula model, 126
stateful inspection firewalls, 210, 330
states
data, 109
system, 124–126
static application security testing (SAST), 410
static firewalls, 210, 329
steganography, 112
STIX (Structured Threat Information eXpression) threat model, 301
storage
assets, 91–92
backup. *See* backup storage
evidence, 176
media, 316
storage area networks (SANs), 350
storage keys in TPM, 136
storage root keys (SRKs) in TPM, 136
strategic goals, organizational, 17
strategies for security evaluations, 261–262
stream ciphers, 152
STRIDE (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege) threat model, 72, 301
strong authentication, 231
strong star property rule in Bell-LaPadula model, 126
structural threats, 71
Structured Threat Information eXpression (STIX) threat model, 301
subjects
certificates, 157
data, 101–102
entities, 10
IAM, 226
security models, 122
subordinate CAs, 157
substitution ciphers, 149
supervisor responsibilities, 19
supervisory control and data acquisition (SCADA) systems, 142–143
supplicants in IEEE 802.1X, 192
supply chain risk management (SCRM)
description, 74
hardware, 75
minimum security requirements, 77
review and questions, 77–79
service level requirements, 77
services, 75–76
software, 75
third-party assessment and monitoring, 76
suppression, fire, 177–178
surges, power, 180
surveillance
CPTED, 174
perimeter security controls, 381
sustainment phase in asset life cycle, 106
switches, 209
symmetric encryption, 151–152
synchronization, code, 408
synthetic transactions, 268
system criticality, 340
system-high security mode, 125
system owner responsibilities, 19
system/security administrators responsibilities, 19
system states in security models, 124–126
systems, description, 9–10, 227

T

tables in database systems, 141–142
tabletop tests
business continuity plans, 375
disaster recovery plans, 368–369
TACACS, TACACS+ (Terminal Access Controller Access Control System), 221, 256
tactical goals, 17
tangible assets, 57, 97
targeted penetration tests, 267
targets in iSCSI, 195
TAXII (Trusted Automated eXchange of Indicator Information) threat model, 301

TCP/IP stack, 186–188
TDMA (time division multiple access), 204
technical controls, 65
technical vulnerabilities, 339
TEEs (trusted execution environments), 137
temperature in intrusion detection systems, 384
Temporal Key Integrity Protocol (TKIP), 200
temporary privilege escalation, 248–249
Ten Commandments of Computer Ethics, 5
tenets of information security
    auditing and accountability, 12
    authentication, 11–12
    authenticity, 12
    authorization, 12
    identification, 11
    nonrepudiation, 12–13
    review and questions, 14–16
    security concepts, 13–14
Terminal Access Controller Access Control System (TACACS, TACACS+), 221, 256
terminations, personnel, 52
territorial reinforcement in CPTED, 174
test coverage analysis in security control testing, 269
tests
    application security, 408–411
    asset life cycle, 106
    business continuity plans, 375–376
    change approval, 345
    definition, 260–261
    disaster recovery plans, 367–372
    patches, 341–342
    security control, 264–272
    software development life cycle, 395
TGS (ticket granting service) in Kerberos, 255
TGTs (ticket granting tickets) in Kerberos, 254–255
third generation programming languages, 404
third parties
    agreements and controls, 52–53
    assessment and monitoring, 76
    communications channels connectivity, 222–223
    identity services, 237–239
    security auditors, 283–284
    security evaluation assessors, 262–263
    security services, 332–333
    software, 417–418
threat hunting, 300
threat intelligence, 298–300
threat modeling
    actors, 58–59, 71–72
    characteristics, 71
    components, 70–71
    methodologies, 72, 300–301
    overview, 70
    review and questions, 73–74
    secure design, 116
    site and facility design, 168–169
threat-vulnerability pairs, 58
threats
    description, 58
    identification, 59
3DES algorithm, 153
ticket granting service (TGS) in Kerberos, 255
ticket granting tickets (TGTs) in Kerberos, 254–255
tickets in Kerberos, 254–255
tight coupling in programming, 405
time division multiple access (TDMA), 204
time-of-check to time-of-use (TOC/TOU) attacks, 138, 164
timeliness in threat intelligence, 299
timing cryptanalytic attacks, 164
TLS (Transport Layer Security), 189–190
TOC/TOU (time-of-check to time-of-use) attacks, 138, 164
tolerance, risk, 60
tool sets in software development, 406
TPMs (Trusted Platform Modules), 135–136
traceroute utility, 188
trade secrets, 30–31
trademarks, 30–31
training
    disaster recovery plans, 364
    personnel security, 389
    security awareness, 80–84
    security process data, 274
transactions
    Clark-Wilson model, 127–128
    synthetic, 268
transborder data flow, 32
transfers
    assets, 93
    personnel, 51–52
transmission media in networks, 212–213
transmission of assets, 92
Transport Layer Security (TLS), 189–190
transportation
    assets, 92
    media, 316
transposition ciphers, 149
travel safety, 388–389
Trike threat model, 72, 301
trust but verify principle
    secure design, 119–120
    site and facility design, 171
Trusted Automated eXchange of Indicator Information (TAXII) threat model, 301
trusted execution environments (TEEs), 137
Trusted Platform Modules (TPMs), 135–136
tumbler locks, 384
turnstiles, 380
Twofish algorithm, 153
Type I hypervisors, 145
Type II hypervisors, 145

U

U.S. government classification systems, 94
UEBA (user and entity behavior analytics), 301–302
unethical (black-hat) hackers, 266
uninterruptible power supplies, 180
unit tests for application security, 409
updates
    criticality, 340–341
    managing, 340
    patches, 341
urgent changes, 346
use case tests in application security, 409
user accounts
    access review, 249–250
    deprovisioning, 246–247
    registration, 246
    security process data, 274

user and entity behavior analytics (UEBA), 301–302
user stories in Agile methodology, 397
users
    data, 101–102
    responsibilities, 19
utilities in site and facility controls, 177

V

vacations, mandatory, 312
validating evaluations, 261–264
VAST (Visual, Agile, and Simple Threat) modeling, 72, 301
vaulting, electronic, 351
vectors, threat, 300
vendors, agreements and controls with, 52–53
verification, backup, 274
versatile memory in TPM, 136
version control
    software, 401
    software configuration management, 408
vertical enactments, 32
very high-level programming languages, 404
Virtual eXtensible Local Area Network (VxLAN), 196, 222
virtual LANs (VLANs), 208, 222
virtual private networks (VPNs), 219–220
virtual storage area networks (VSANs), 222
virtualized networks communications channels, 222
virtualized system vulnerabilities, 145
Visual, Agile, and Simple Threat (VAST) modeling, 72, 301
VLANs (virtual LANs), 208, 222
voice communications, 215–218
Voice over Internet Protocol (VoIP), 195, 217
VPNs (virtual private networks), 219–220
VSANs (virtual storage area networks), 222
vulnerabilities
    assessments, 59, 265
    client-based systems, 140
    cloud-based systems, 144–145
    containerization, 146
    cryptographic systems, 142
    database systems, 141–142
    description, 57–58
    distributed systems, 141
    edge computing systems, 146–147
    embedded systems, 143–144
    high-performance computing systems, 146
    identifying, 59–60
    industrial control systems, 142–143
    Internet of Things, 143
    microservices, 146
    review and questions, 147–148
    server-based systems, 140–141
    serverless functions, 146
    source-code level, 420–421
    virtualized systems, 145
vulnerability management
    nontechnical, 340
    review and questions, 342–344
    technical, 339
vulnerability testing in application security, 410
VxLAN (Virtual eXtensible Local Area Network), 196, 222

W

WAFs (web application firewalls), 330
walk-through tests
    business continuity plans, 375
    disaster recovery plans, 369
warded locks, 383
warm sites, 353
Wassenaar Arrangement, 31
water sprinkler systems, 179
Waterfall methodology, 395–396
wave pattern motion detectors, 385
weaknesses in source-code level, 420–421
web application firewalls (WAFs), 330
well-formed transactions in Clark-Wilson model, 127–128
WEP (Wired Equivalent Privacy), 200
wet pipe water sprinkler systems, 179
white-box penetration tests, 267
white-hat (ethical) hackers, 266
whitelisting, 327–328
Wi-Fi
    fundamentals, 199
    overview, 199
    physical standards, 199–201
    security, 200–202
Wi-Fi Protected Access (WPA), 200
Wired Equivalent Privacy (WEP), 200
wireless technologies
    Bluetooth, 202
    cellular networks, 204–205
    introduction, 197
    Li-Fi, 203–204
    satellites, 203
    theory and signaling, 197–198
    Wi-Fi, 199–202
    Zigbee, 202–203
wiring closets, 175
work area security, 176
work functions in cryptosystems, 151
workplace ethics statements and policies, 4–5
WPA (Wi-Fi Protected Access), 200

X

XP (Extreme Programming), 398

Z

zero-knowledge penetration tests, 267
zero trust principle
    secure design, 119
    site and facility design, 171
Zigbee technology, 202–203
zones in perimeter security, 378–379