PEARSON IT CERTIFICATION

Save 10% on Exam Voucher
See Inside



Practice Tests

Flash Cards

Review Exercises

Study Planner

Cert Guide

Advance your IT career with hands-on learning

CompTIA®

Security+

SY0-601

OMAR SANTOS
RON TAYLOR
JOSEPH MLODZIANOWSKI

Special Offer

Save 80% on Premium Edition eBook and Practice Test

The *CompTIA Security+ SY0-601 Cert Guide Premium Edition eBook and Practice Test* provides three eBook files (PDF, EPUB, and MOBI/Kindle) to read on your preferred device and an enhanced edition of the Pearson Test Prep practice test software. You also receive two additional practice exams with links for every question mapped to the PDF eBook.

See the card insert in the back of the book for your Pearson Test Prep activation code and special offers.

CompTIA® Security+ SY0-601 Cert Guide

Omar Santos
Ron Taylor
Joseph Mlodzianowski



CompTIA® Security+ SY0-601 Cert Guide

Copyright © 2022 by Pearson Education, Inc.

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

ISBN-13: 978-0-13-677031-2

ISBN-10: 0-13-677031-2

Library of Congress Control Number: 2021935686

ScoutAutomatedPrintCode

Trademarks

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

Warning and Disclaimer

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

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Editor-in-Chief
Mark Taub

Product Line Manager
Brett Bartow

Executive Editor
Nancy Davis

Development Editor
Christopher A. Cleveland

Managing Editor
Sandra Schroeder

Senior Project Editor
Tonya Simpson

Copy Editor
Chuck Hutchinson

Indexer
Erika Millen

Proofreader
Abigail Manheim

Technical Editor
Chris Crayton

Publishing Coordinator
Cindy Teeters

Cover Designer
Chuti Prasertsith

Compositor
codeMantra

Contents at a Glance

Introduction xliv

Part I: Threats, Attacks, and Vulnerabilities

CHAPTER 1 Comparing and Contrasting Different Types of Social Engineering Techniques 3

CHAPTER 2 Analyzing Potential Indicators to Determine the Type of Attack 29

CHAPTER 3 Analyzing Potential Indicators Associated with Application Attacks 61

CHAPTER 4 Analyzing Potential Indicators Associated with Network Attacks 95

CHAPTER 5 Understanding Different Threat Actors, Vectors, and Intelligence Sources 117

CHAPTER 6 Understanding the Security Concerns Associated with Various Types of Vulnerabilities 133

CHAPTER 7 Summarizing the Techniques Used in Security Assessments 171

CHAPTER 8 Understanding the Techniques Used in Penetration Testing 193

Part II: Architecture and Design

CHAPTER 9 Understanding the Importance of Security Concepts in an Enterprise Environment 209

CHAPTER 10 Summarizing Virtualization and Cloud Computing Concepts 227

CHAPTER 11 Summarizing Secure Application Development, Deployment, and Automation Concepts 253

CHAPTER 12 Summarizing Authentication and Authorization Design Concepts 285

CHAPTER 13 Implementing Cybersecurity Resilience 311

CHAPTER 14 Understanding the Security Implications of Embedded and Specialized Systems 335

CHAPTER 15 Understanding the Importance of Physical Security Controls 367

CHAPTER 16 Summarizing the Basics of Cryptographic Concepts 391

Part III: Implementation

CHAPTER 17 Implementing Secure Protocols 423

CHAPTER 18 Implementing Host or Application Security Solutions 447

CHAPTER 19 Implementing Secure Network Designs 483

CHAPTER 20 Installing and Configuring Wireless Security Settings 547

CHAPTER 21 Implementing Secure Mobile Solutions 567

CHAPTER 22 Applying Cybersecurity Solutions to the Cloud 595

CHAPTER 23 Implementing Identity and Account Management Controls 619

CHAPTER 24 Implementing Authentication and Authorization Solutions 651

CHAPTER 25 Implementing Public Key Infrastructure 685

Part IV: Operations and Incident Response

CHAPTER 26 Using the Appropriate Tool to Assess Organizational Security 703

CHAPTER 27 Summarizing the Importance of Policies, Processes, and Procedures for Incident Response 755

CHAPTER 28 Using Appropriate Data Sources to Support an Investigation 781

CHAPTER 29 Applying Mitigation Techniques or Controls to Secure an Environment 819

CHAPTER 30 Understanding the Key Aspects of Digital Forensics 837

Part V: Governance, Risk, and Compliance

CHAPTER 31 Comparing and Contrasting the Various Types of Controls 865

CHAPTER 32 Understanding the Importance of Applicable Regulations, Standards, or Frameworks That Impact Organizational Security Posture 875

CHAPTER 33 Understanding the Importance of Policies to Organizational Security 893

CHAPTER 34 Summarizing Risk Management Processes and Concepts 913

CHAPTER 35 Understanding Privacy and Sensitive Data Concepts in Relation to Security 935

Part VI: Final Preparation

CHAPTER 36 Final Preparation 953

Glossary of Key Terms 955

APPENDIX A Answers to the "Do I Know This Already?" Quizzes and Review Questions 1023

APPENDIX B *CompTIA Security+ (SY0-601) Cert Guide* Exam Updates 1087

Index 1089

Online Elements:

APPENDIX C Study Planner

Glossary of Key Terms

Table of Contents

Introduction xliv

Part I: Threats, Attacks, and Vulnerabilities

Chapter 1 Comparing and Contrasting Different Types of Social Engineering Techniques 3

"Do I Know This Already?" Quiz 3

Foundation Topics 7

Social Engineering Fundamentals 7

Phishing and Spear Phishing 9

Smishing 12

Vishing 12

Spam and Spam over Internet Messaging (SPIM) 13

Dumpster Diving 13

Shoulder Surfing 14

Pharming 14

Piggybacking or Tailgating 15

Eliciting Information 15

Whaling 16

Prepending 17

Identity Fraud 17

Invoice Scams 17

Credential Harvesting 18

Reconnaissance 18

Hoaxes 19

Impersonation or Pretexting 19

Eavesdropping 19

Baiting 20

Watering Hole Attack 20

Typo Squatting 20

Influence Campaigns, Principles of Social Engineering, and Reasons for Effectiveness 21

User Security Awareness Education 22
Chapter Review Activities 24
Review Key Topics 24
Define Key Terms 25
Review Questions 26
Chapter 2 Analyzing Potential Indicators to Determine the Type of Attack 29
"Do I Know This Already?" Quiz 29
Foundation Topics 33
Malicious Software (Malware) 33
Ransomware and Cryptomalware 33
Trojans 35
Remote Access Trojans (RATs) and Rootkits 35
Worms 36
Fileless Virus 37
Command and Control, Bots, and Botnets 37
Logic Bombs 39
Potentially Unwanted Programs (PUPs) and Spyware 40
Keyloggers 42
Backdoors 43
Malware Delivery Mechanisms 43
You Can't Save Every Computer from Malware! 45
Password Attacks 45
Dictionary-based and Brute-force Attacks 45
Password Spraying 46
Offline and Online Password Cracking 46
Rainbow Tables 47
Plaintext/Unencrypted 47
Physical Attacks 48
Malicious Flash Drives 48
Malicious Universal Serial Bus (USB) Cables 48
Card Cloning Attacks 48
Skimming 49

Adversarial Artificial Intelligence 50
Tainted Training Data for Machine Learning 50
Security of Machine Learning Algorithms 50
Supply-Chain Attacks 51
Cloud-based vs. On-premises Attacks 52
Cloud Security Threats 52
Cloud Computing Attacks 54
Cryptographic Attacks 55
Collision 55
Birthday 56
Downgrade 56
Chapter Review Activities 57
Review Key Topics 57
Define Key Terms 58
Review Questions 59

Chapter 3 Analyzing Potential Indicators Associated with Application Attacks 61

"Do I Know This Already?" Quiz 61
Foundation Topics 67
Privilege Escalation 67
Cross-Site Scripting (XSS) Attacks 68
Injection Attacks 70
Structured Query Language (SQL) Injection Attacks 70
SQL Injection Categories 73
Dynamic Link Library (DLL) Injection Attacks 74
Lightweight Directory Access Protocol (LDAP) Injection Attacks 74
Extensible Markup Language (XML) Injection Attacks 74
Pointer/Object Dereference 75
Directory Traversal 76
Buffer Overflows 77
Arbitrary Code Execution/Remote Code Execution 78
Race Conditions 79
Error Handling 79

Improper Input Handling 80
Compile-Time Errors vs. Runtime Errors 81
Replay Attacks 82
Request Forgeries 85
Application Programming Interface (API) Attacks 86
Resource Exhaustion 87
Memory Leaks 88
Secure Socket Layer (SSL) Stripping 88
Driver Manipulation 89
Pass the Hash 89
Chapter Review Activities 90
Review Key Topics 90
Define Key Terms 92
Review Questions 92
Chapter 4 Analyzing Potential Indicators Associated with Network Attacks 95
"Do I Know This Already?" Quiz 95
Foundation Topics 98
Wireless Attacks 98
Evil Twin Attacks 98
Rogue Access Points 99
Bluesnarfing Attacks 99
Bluejacking Attacks 100
Disassociation and Deauthentication Attacks 101
Jamming Attacks 102
Radio Frequency Identifier (RFID) Attacks 102
Near-Field Communication (NFC) Attacks 102
Initialization Vector (IV) Attacks 103
On-Path Attacks 103
Layer 2 Attacks 105
Address Resolution Protocol (ARP) Poisoning Attacks 105
Media Access Control (MAC) Flooding Attacks 106
MAC Cloning Attacks 106
Best Practices to Protect Against Layer 2 Attacks 106

Domain Name System (DNS) Attacks 107
Domain Hijacking Attacks 108
DNS Poisoning Attacks 108
Uniform Resource Locator (URL) Redirection Attacks 110
Domain Reputation 110
Distributed Denial-of-Service (DDoS) Attacks 111
Malicious Code or Script Execution Attacks 113
Chapter Review Activities 114
Review Key Topics 114
Define Key Terms 115
Review Questions 115

Chapter 5 Understanding Different Threat Actors, Vectors, and Intelligence Sources 117

"Do I Know This Already?" Quiz 117
Foundation Topics 120
Actors and Threats 120
Attributes of Threat Actors 122
Attack Vectors 122
Threat Intelligence and Threat Intelligence Sources 123
Structured Threat Information eXpression (STIX) and the Trusted Automated eXchange of Indicator Information (TAXII) 125
Research Sources 127
The MITRE ATT&CK Framework 128
Chapter Review Activities 129
Review Key Topics 129
Define Key Terms 130
Review Questions 131

Chapter 6 Understanding the Security Concerns Associated with Various Types of Vulnerabilities 133

"Do I Know This Already?" Quiz 133
Foundation Topics 137
Cloud-based vs. On-premises Vulnerabilities 137
Other "Cloud"-based Concerns 143
Server Defense 144

File Servers 144
Network Controllers 144
Email Servers 145
Web Servers 146
FTP Server 147
Zero-day Vulnerabilities 149
Weak Configurations 150
Third-party Risks 155
Improper or Weak Patch Management 160
Patches and Hotfixes 161
Patch Management 163
Legacy Platforms 165
The Impact of Cybersecurity Attacks and Breaches 165
Chapter Review Activities 166
Review Key Topics 166
Define Key Terms 167
Review Questions 168
Chapter 7 Summarizing the Techniques Used in Security Assessments 171
"Do I Know This Already?" Quiz 171
Foundation Topics 175
Threat Hunting 175
Security Advisories and Bulletins 177
Vulnerability Scans 180
Credentialed vs. Noncredentialed 182
Intrusive vs. Nonintrusive 182
Common Vulnerability Scoring System (CVSS) 182
Logs and Security Information and Event Management (SIEM) 186
Security Orchestration, Automation, and Response (SOAR) 188
Chapter Review Activities 189
Review Key Topics 189
Define Key Terms 190
Review Questions 190

Chapter 8 Understanding the Techniques Used in Penetration Testing 193

"Do I Know This Already?" Quiz 193

Foundation Topics 197

Penetration Testing 197

Bug Bounties vs. Penetration Testing 202

Passive and Active Reconnaissance 203

Exercise Types 205

Chapter Review Activities 206

Review Key Topics 206

Define Key Terms 207

Review Questions 207

Part II: Architecture and Design

Chapter 9 Understanding the Importance of Security Concepts in an Enterprise Environment 209

"Do I Know This Already?" Quiz 209

Foundation Topics 213

Configuration Management 213

Data Sovereignty and Data Protection 214

Secure Sockets Layer (SSL)/Transport Layer Security (TLS) Inspection 215

API Considerations 216

Data Masking and Obfuscation 216

Encryption at Rest, in Transit/Motion, and in Processing 218

Hashing 218

Rights Management 219

Geographical Considerations 220

Data Breach Response and Recovery Controls 220

Site Resiliency 221

Deception and Disruption 222

Fake Telemetry 223

DNS Sinkhole 223

Chapter Review Activities 224

Review Key Topics 224
Define Key Terms 225
Review Questions 225

Chapter 10 Summarizing Virtualization and Cloud Computing Concepts 227

"Do I Know This Already?" Quiz 227
Foundation Topics 231
Cloud Models 231
Public, Private, Hybrid, and Community Clouds 232
Cloud Service Providers 233
Cloud Architecture Components 234
Fog and Edge Computing 234
Thin Clients 235
Containers 236
Microservices and APIs 240
Infrastructure as Code 241
Serverless Architecture 243
Services Integration 246
Resource Policies 246
Transit Gateway 246
Virtual Machine (VM) Sprawl Avoidance and VM Escape Protection 247
Understanding and Avoiding VM Sprawl 247
Protecting Against VM Escape Attacks 248
Chapter Review Activities 250
Review Key Topics 250
Define Key Terms 251
Review Questions 251

Chapter 11 Summarizing Secure Application Development, Deployment, and Automation Concepts 253

"Do I Know This Already?" Quiz 253
Foundation Topics 257
Software Development Environments and Methodologies 257
Application Provisioning and Deprovisioning 260
Software Integrity Measurement 261

Secure Coding Techniques 261
Core SDLC and DevOps Principles 263
Programming Testing Methods 266
Compile-Time Errors vs. Runtime Errors 266
Input Validation 267
Static and Dynamic Code Analysis 269
Fuzz Testing 269
Programming Vulnerabilities and Attacks 270
Testing for Backdoors 271
Memory/Buffer Vulnerabilities 271
XSS and XSRF 272
More Code Injection Examples 273
Directory Traversal 274
Zero-Day Attack 275
Open Web Application Security Project (OWASP) 276
Software Diversity 278
Automation/Scripting 278
Elasticity and Scalability 279
Chapter Review Activities 280
Review Key Topics 280
Define Key Terms 281
Review Questions 281

Chapter 12 Summarizing Authentication and Authorization Design Concepts 285

"Do I Know This Already?" Quiz 285
Foundation Topics 289
Authentication Methods 289
Directory Services 291
Federations 292
Attestation 294
Authentication Methods and Technologies 295
Time-Based One-Time Password (TOTP) 295
HMAC-Based One-Time Password (HOTP) 295

Short Message Service (SMS) 296
Token Key 297
Static Codes 298
Authentication Applications 298
Push Notifications 299
Phone Call Authentication 299
Smart Card Authentication 300
Biometrics 300
Fingerprints 300
Retina 301
Iris 301
Facial 301
Voice 302
Vein 302
Gait Analysis 302
Efficacy Rates 302
False Acceptance 303
False Rejection 303
Crossover Error Rate 304
Multifactor Authentication (MFA) Factors and Attributes 304
Authentication, Authorization, and Accounting (AAA) 306
Cloud vs. On-premises Requirements 306
Chapter Review Activities 308
Review Key Topics 308
Define Key Terms 308
Review Questions 308
Chapter 13 Implementing Cybersecurity Resilience 311
"Do I Know This Already?" Quiz 311
Foundation Topics 315
Redundancy 315
Geographic Dispersal 315
Disk Redundancy 315
Redundant Array of Inexpensive Disks 316

Multipath 319
Network Resilience 319
Load Balancers 319
Network Interface Card (NIC) Teaming 320
Power Resilience 320
Uninterruptible Power Supply (UPS) 320
Generators 321
Dual Supply 321
Managed Power Distribution Units (PDUs) 322
Replication 323
Storage Area Network 323
Virtual Machines 324
On-premises vs. Cloud 325
Backup Types 326
Full Backup 328
Differential Backup 328
Incremental Backup 328
Non-persistence 328
High Availability 329
Restoration Order 330
Diversity 331
Technologies 331
Vendors 331
Crypto 331
Controls 332
Chapter Review Activities 332
Review Key Topics 332
Define Key Terms 333
Review Questions 333
Chapter 14 Understanding the Security Implications of Embedded and Specialized Systems 335
"Do I Know This Already?" Quiz 335
Foundation Topics 339

Embedded Systems 339
Supervisory Control and Data Acquisition (SCADA)/Industrial Control Systems (ICS) 341
Internet of Things (IoT) 344
Specialized Systems 346
Medical Systems 347
Vehicles 347
Aircraft 348
Smart Meters 350
Voice over IP (VoIP) 351
Heating, Ventilation, and Air Conditioning (HVAC) 352
Drones 353
Multifunction Printers (MFP) 354
Real-Time Operating Systems (RTOS) 355
Surveillance Systems 355
System on a Chip (SoC) 356
Communication Considerations 357
5G 357
NarrowBand 358
Baseband Radio 359
Subscriber Identity Module (SIM) Cards 360
Zigbee 360
Embedded System Constraints 361
Power 361
Compute 361
Network 362
Crypto 362
Inability to Patch 362
Authentication 363
Range 363
Cost 363
Implied Trust 363

Chapter Review Activities 364
Review Key Topics 364
Define Key Terms 365
Review Questions 365
Chapter 15 Understanding the Importance of Physical Security Controls 367
"Do I Know This Already?" Quiz 367
Foundation Topics 370
Bollards/Barricades 370
Access Control Vestibules 372
Badges 373
Alarms 374
Signage 374
Cameras 375
Closed-Circuit Television (CCTV) 376
Industrial Camouflage 377
Personnel 377
Locks 378
USB Data Blockers 379
Lighting 380
Fencing 380
Fire Suppression 381
Sensors 381
Drones 382
Visitor Logs 383
Faraday Cages 383
Air Gap 384
Screened Subnet (Previously Known as Demilitarized Zone [DMZ]) 384
Protected Cable Distribution 385
Secure Areas 385
Secure Data Destruction 386
Chapter Review Activities 387
Review Key Topics 387
Define Key Terms 388
Review Questions 389

Chapter 16 Summarizing the Basics of Cryptographic Concepts 391

"Do I Know This Already?" Quiz 391

Foundation Topics 395

Digital Signatures 395

Key Length 396

Key Stretching 397

Salting 397

Hashing 398

Key Exchange 399

Elliptic-Curve Cryptography 399

Perfect Forward Secrecy 400

Quantum 401

- Communications 401
- Computing 402

Post-Quantum 402

Ephemeral 403

Modes of Operation 403

- Electronic Code Book Mode 404
- Cipher Block Chaining Mode 405
- Cipher Feedback Mode 406
- Output Feedback Mode 407
- Counter Mode 408

Blockchain 409

Cipher Suites 410

Symmetric vs. Asymmetric Encryption 411

Lightweight Cryptography 414

Steganography 415

- Audio Steganography 415
- Video Steganography 416
- Image Steganography 416

Homomorphic Encryption 417

Common Use Cases 417

Limitations 418

Chapter Review Activities 420
Review Key Topics 420
Define Key Terms 421
Review Questions 421

Part III: Implementation

Chapter 17 Implementing Secure Protocols 423

"Do I Know This Already?" Quiz 423
Foundation Topics 426
Protocols 426
Domain Name System Security Extensions 426
SSH 427
Secure/Multipurpose Internet Mail Extensions 428
Secure Real-Time Transport Protocol 430
Lightweight Directory Access Protocol over SSL 432
File Transfer Protocol, Secure 432
Secure (or SSH) File Transfer Protocol 434
Simple Network Management Protocol Version 3 434
Hypertext Transfer Protocol over SSL/TLS 436
IPsec 437
Authentication Header/Encapsulating Security Payloads 437
Tunnel/Transport 438
Post Office Protocol/Internet Message Access Protocol 438
Use Cases 439
Voice and Video 440
Time Synchronization 440
Email and Web 441
File Transfer 441
Directory Services 442
Remote Access 442
Domain Name Resolution 442
Routing and Switching 443
Network Address Allocation 443
Subscription Services 444

Chapter Review Activities 444
Review Key Topics 444
Define Key Terms 445
Review Questions 445
Chapter 18 Implementing Host or Application Security Solutions 447
"Do I Know This Already?" Quiz 447
Foundation Topics 451
Endpoint Protection 451
Antivirus 451
Antimalware 452
Endpoint Detection and Response 452
Data Loss Prevention 453
Next-Generation Firewall 453
Host-based Intrusion Prevention System 454
Host-based Intrusion Detection System 456
Host-based Firewall 457
Boot Integrity 458
Boot Security/Unified Extensible Firmware Interface 459
Measured Boot 459
Boot Attestation 460
Database 461
Tokenization 461
Salting 462
Hashing 463
Application Security 463
Input Validations 464
Secure Cookies 465
Hypertext Transfer Protocol Headers 465
End-to-End Headers 466
Hop-by-Hop Headers 466
Code Signing 466
Allow List 467
Block List/Deny List 467

Secure Coding Practices 468
Static Code Analysis 468
Manual Code Review 470
Dynamic Code Analysis 470
Fuzzing 471
Hardening 471
Open Ports and Services 471
Registry 472
Disk Encryption 473
Operating System 473
Patch Management 474
Self-Encrypting Drive/Full-Disk Encryption 475
OPAL 476
Hardware Root of Trust 476
Trusted Platform Module 477
Sandboxing 478
Chapter Review Activities 479
Review Key Topics 479
Define Key Terms 481
Review Questions 481
Chapter 19 Implementing Secure Network Designs 483
"Do I Know This Already?" Quiz 483
Foundation Topics 488
Load Balancing 488
Active/Active 488
Active/Passive 488
Scheduling 488
Virtual IP 488
Persistence 489
Network Segmentation 489
Application-Based Segmentation and Microsegmentation 489
Virtual Local Area Network 490
Screened Subnet 491

East-West Traffic 492
Intranets and Extranets 492
Zero Trust 494
Virtual Private Network 494
Remote Access vs. Site-to-Site 496
IPsec 497
IKEv1 Phase 1 498
IKEv1 Phase 2 501
IKEv2 504
SSL/TLS 505
HTML5 508
Layer 2 Tunneling Protocol 508
DNS 509
Network Access Control 510
Out-of-Band Management 510
Port Security 511
Broadcast Storm Prevention 512
Bridge Protocol Data Unit Guard 512
Loop Prevention 512
Dynamic Host Configuration Protocol Snooping 512
Media Access Control Filtering 513
Network Appliances 513
Jump Servers 514
Proxy Servers 514
Network-Based Intrusion Detection System/Network-Based Intrusion Prevention System 516
NIDS 517
NIPS 518
Summary of NIDS vs. NIPS 519
Signature-Based 520
Heuristic/Behavior 521
Anomaly 521
Inline vs. Passive 523

HSM 524
Sensors 524
Collectors 525
Aggregators 526
Firewalls 526
Hardware vs. Software 534
Appliance vs. Host-based vs. Virtual 534
Access Control List 535
Route Security 535
Quality of Service 536
Implications of IPv6 536
Port Spanning/Port Mirroring 537
Monitoring Services 538
Performance Baselining 539
File Integrity Monitors 542
Chapter Review Activities 542
Review Key Topics 542
Define Key Terms 543
Review Questions 544
Chapter 20 Installing and Configuring Wireless Security Settings 547
"Do I Know This Already?" Quiz 547
Foundation Topics 551
Cryptographic Protocols 551
Wi-Fi Protected Access 2 (WPA2) 551
Wi-Fi Protected Access 3 (WPA3) 551
Counter-mode/CBC-MAC Protocol (CCMP) 552
Simultaneous Authentication of Equals 552
Wireless Cryptographic Protocol Summary 552
Authentication Protocols 553
802.1X and EAP 553
IEEE 802.1x 556
Remote Authentication Dial-In User Service (RADIUS) Federation 556

Methods 557
Wi-Fi Protected Setup 558
Captive Portals 559
Installation Considerations 559
Controller and Access Point Security 562
Wireless Access Point Vulnerabilities 563
Chapter Review Activities 564
Review Key Topics 564
Define Key Terms 564
Review Questions 565
Chapter 21 Implementing Secure Mobile Solutions 567
"Do I Know This Already?" Quiz 567
Foundation Topics 570
Connection Methods and Receivers 570
RFID and NFC 571
More Wireless Connection Methods and Receivers 572
Secure Implementation Best Practices 573
Mobile Device Management 574
MDM Security Feature Concerns: Application and Content Management 576
MDM Security Feature Concerns: Remote Wipe, Geofencing, Geolocation, Screen Locks, Passwords and PINs, Full Device Encryption 578
Mobile Device Management Enforcement and Monitoring 581
Mobile Devices 585
MDM/Unified Endpoint Management 587
SEAndroid 588
Deployment Models 588
Secure Implementation of BYOD, CYOD, and COPE 589
Chapter Review Activities 591
Review Key Topics 591
Define Key Terms 592
Review Questions 592

Chapter 22 Applying Cybersecurity Solutions to the Cloud 595

"Do I Know This Already?" Quiz 595

Foundation Topics 598

Cloud Security Controls 598

Security Assessment in the Cloud 598

Understanding the Different Cloud Security Threats 598

Cloud Computing Attacks 601

High Availability Across Zones 603

Resource Policies 603

Integration and Auditing 604

Secrets Management 604

Storage 605

Permissions 605

Encryption 605

Replication 605

High Availability 606

Network 606

Virtual Networks 606

Public and Private Subnets 606

Segmentation 607

API Inspection and Integration 607

Compute 607

Security Groups 607

Dynamic Resource Allocation 607

Instance Awareness 608

Virtual Private Cloud Endpoint 608

Container Security 608

Summary of Cloud Security Controls 609

Solutions 611

CASB 611

Application Security 612

Next-Generation Secure Web Gateway 613

Firewall Considerations in a Cloud Environment 613

Cost 613
Need for Segmentation 613
Open Systems Interconnection Layers 614
Summary of Cybersecurity Solutions to the Cloud 614
Cloud Native Controls vs. Third-Party Solutions 615
Chapter Review Activities 615
Review Key Topics 615
Define Key Terms 616
Review Questions 616
Chapter 23 Implementing Identity and Account Management Controls 619
"Do I Know This Already?" Quiz 619
Foundation Topics 623
Identity 623
Identity Provider (IdP) 623
Authentication 625
Authentication by Knowledge 625
Authentication by Ownership 625
Authentication by Characteristic Attributes 625
Certificates 626
Tokens 627
SSH Keys 628
Smart Cards 629
Account Types 629
Account Policies 633
Introduction to Identity and Access Management 633
Phases of the Identity and Access Lifecycle 633
Registration and Identity Validation 634
Privileges Provisioning 635
Access Review 635
Access Revocation 635
Password Management 636
Password Creation 636
Attribute-Based Access Control (ABAC) 638

Rights, Permissions, and Policies 640
Users, Groups, and Account Permissions 640
Permission Inheritance and Propagation 645
Chapter Review Activities 647
Review Key Topics 647
Define Key Terms 647
Review Questions 648
Chapter 24 Implementing Authentication and Authorization Solutions 651
"Do I Know This Already?" Quiz 651
Foundation Topics 655
Authentication Management 655
Password Keys 655
Password Vaults 655
Trusted Platform Module 656
Hardware Security Modules 656
Knowledge-Based Authentication 656
Authentication/Authorization 657
Security Assertion Markup Language 659
OAuth 661
OpenID and OpenID Connect 663
802.1X and EAP 664
LDAP 667
Kerberos and Mutual Authentication 668
Remote Authentication Technologies 670
Remote Access Service 670
RADIUS versus TACACS+ 672
Access Control Schemes 674
Discretionary Access Control 674
Mandatory Access Control 676
Role-Based Access Control 677
Attribute-Based Access Control 678
Rule-Based Access Control 678
Conditional Access 678
Privileged Access Management 678

Summary of Access Control Models 679
Access Control Wise Practices 680
Chapter Review Activities 681
Review Key Topics 681
Define Key Terms 682
Review Questions 682

Chapter 25 Implementing Public Key Infrastructure 685

"Do I Know This Already?" Quiz 685
Foundation Topics 688
Public Key Infrastructure 688
Key Management 688
Certificate Authorities 689
Certificate Attributes 691
Subject Alternative Name 693
Expiration 693
Types of Certificates 694
SSL Certificate Types 694
Certificate Chaining 696
Certificate Formats 697
PKI Concepts 698
Trust Model 698
Certificate Pinning 698
Stapling, Key Escrow, Certificate Chaining, Online vs. Offline CA 698
Chapter Review Activities 700
Review Key Topics 700
Define Key Terms 700
Review Questions 701

Part IV: Operations and Incident Response

Chapter 26 Using the Appropriate Tool to Assess Organizational Security 703

"Do I Know This Already?" Quiz 703
Foundation Topics 707
Network Reconnaissance and Discovery 707
tracert/traceroute 707

nslookup/dig 709
ipconfig/ifconfig 710
nmap 711
ping/pathping 714
hping 717
netstat 718
netcat 720
IP Scanners 721
arp 721
route 723
curl 724
theHarvester 725
sn1per 726
scanless 727
dnsenum 728
Nessus 730
Cuckoo 731
File Manipulation 732
head 733
tail 734
cat 734
grep 735
chmod 736
Logger 737
Shell and Script Environments 738
SSH 739
PowerShell 740
Python 741
OpenSSL 741
Packet Capture and Replay 742
Tcpreplay 742
Tcpdump 742
Wireshark 743

Forensics 744
dd 744
Memdump 745
WinHex 746
FTK Imager 747
Autopsy 747
Exploitation Frameworks 747
Password Crackers 748
Data Sanitization 750
Chapter Review Activities 750
Review Key Topics 750
Define Key Terms 752
Review Questions 752

Chapter 27 Summarizing the Importance of Policies, Processes, and Procedures for Incident Response 755

"Do I Know This Already?" Quiz 755
Foundation Topics 760
Incident Response Plans 760
Incident Response Process 761
Preparation 762
Identification 763
Containment 763
Eradication 764
Recovery 764
Lessons Learned 764
Exercises 765
Tabletop 765
Walkthroughs 766
Simulations 766
Attack Frameworks 767
MITRE ATT&CK 767
The Diamond Model of Intrusion Analysis 768
Cyber Kill Chain 770

Stakeholder Management 771
Communication Plan 771
Disaster Recovery Plan 772
Business Continuity Plan 773
Continuity of Operations Planning (COOP) 774
Incident Response Team 775
Retention Policies 776
Chapter Review Activities 776
Review Key Topics 776
Define Key Terms 777
Review Questions 778

Chapter 28 Using Appropriate Data Sources to Support an Investigation 781

"Do I Know This Already?" Quiz 781
Foundation Topics 785
Vulnerability Scan Output 785
SIEM Dashboards 786
Sensors 787
Sensitivity 788
Trends 788
Alerts 788
Correlation 788
Log Files 789
Network 790
System 791
Application 792
Security 793
Web 794
DNS 795
Authentication 796
Dump Files 797
VoIP and Call Managers 799
Session Initiation Protocol Traffic 800

syslog/rsyslog/syslog-ng 800
journalctl 802
NXLog 803
Bandwidth Monitors 804
Metadata 805
Email 808
Mobile 808
Web 808
File 809
NetFlow/sFlow 809
NetFlow 809
sFlow 810
IPFIX 811
Protocol Analyzer Output 813
Chapter Review Activities 814
Review Key Topics 814
Define Key Terms 816
Review Questions 816

Chapter 29 Applying Mitigation Techniques or Controls to Secure an Environment 819

"Do I Know This Already?" Quiz 819
Foundation Topics 822
Reconfigure Endpoint Security Solutions 822
Application Approved Lists 822
Application Block List/Deny List 822
Quarantine 823
Configuration Changes 824
Firewall Rules 825
MDM 825
Data Loss Prevention 828
Content Filter/URL Filter 828
Update or Revoke Certificates 829
Isolation 830

Containment 830
Segmentation 831
SOAR 832
Runbooks 833
Playbooks 834
Chapter Review Activities 834
Review Key Topics 834
Define Key Terms 835
Review Questions 835
Chapter 30 Understanding the Key Aspects of Digital Forensics 837
"Do I Know This Already?" Quiz 837
Foundation Topics 842
Documentation/Evidence 842
Legal Hold 842
Video 842
Admissibility 843
Chain of Custody 844
Timelines of Sequence of Events 844
Timestamps 844
Time Offset 845
Tags 845
Reports 846
Event Logs 846
Interviews 846
Acquisition 847
Order of Volatility 848
Disk 848
Random-Access Memory 848
Swap/Pagefile 849
Operating System 850
Device 850
Firmware 851
Snapshot 851
Cache 852

Network 852
Artifacts 853
On-premises vs. Cloud 853
Right-to-Audit Clauses 854
Regulatory/Jurisdiction 855
Data Breach Notification Laws 855
Integrity 856
Hashing 856
Checksums 857
Provenance 857
Preservation 858
E-discovery 858
Data Recovery 859
Nonrepudiation 859
Strategic Intelligence/Counterintelligence 860
Chapter Review Activities 860
Review Key Topics 860
Define Key Terms 862
Review Questions 862

Part V: Governance, Risk, and Compliance

Chapter 31 Comparing and Contrasting the Various Types of Controls 865
"Do I Know This Already?" Quiz 865
Foundation Topics 868
Control Category 868
Managerial Controls 868
Operational Controls 868
Technical Controls 868
Summary of Control Categories 869
Control Types 869
Preventative Controls 869
Detective Controls 869
Corrective Controls 870
Deterrent Controls 870

Compensating Controls 871
Physical Controls 871
Summary of Control Types 872
Chapter Review Activities 873
Review Key Topics 873
Define Key Terms 873
Review Questions 873

Chapter 32 Understanding the Importance of Applicable Regulations, Standards, or Frameworks That Impact Organizational Security Posture 875

"Do I Know This Already?" Quiz 875
Foundation Topics 878
Regulations, Standards, and Legislation 878
General Data Protection Regulation 879
National, Territory, or State Laws 879
Payment Card Industry Data Security Standard (PCI DSS) 881
Key Frameworks 881
Benchmarks and Secure Configuration Guides 885
Security Content Automation Protocol 885
Chapter Review Activities 889
Review Key Topics 889
Define Key Terms 889
Review Questions 890

Chapter 33 Understanding the Importance of Policies to Organizational Security 893

"Do I Know This Already?" Quiz 894
Foundation Topics 897
Personnel Policies 897
Privacy Policies 897
Acceptable Use 898
Separation of Duties/Job Rotation 898
Mandatory Vacations 898
Onboarding and Offboarding 899
Personnel Security Policies 900

Diversity of Training Techniques 900
User Education and Awareness Training 901
Third-Party Risk Management 902
Data Concepts 904
Understanding Classification and Governance 904
Data Retention 906
Credential Policies 906
Organizational Policies 908
Change Management and Change Control 909
Asset Management 909
Chapter Review Activities 910
Review Key Topics 910
Define Key Terms 910
Review Questions 911

Chapter 34 Summarizing Risk Management Processes and Concepts 913

"Do I Know This Already?" Quiz 913
Foundation Topics 917
Risk Types 917
Risk Management Strategies 918
Risk Analysis 919
Qualitative Risk Assessment 921
Quantitative Risk Assessment 922
Disaster Analysis 924
Business Impact Analysis 926
Disaster Recovery Planning 928
Chapter Review Activities 930
Review Key Topics 930
Define Key Terms 931
Review Questions 931

Chapter 35 Understanding Privacy and Sensitive Data Concepts in Relation to Security 935

"Do I Know This Already?" Quiz 935
Foundation Topics 940
Organizational Consequences of Privacy and Data Breaches 940

Notifications of Breaches 941
Data Types and Asset Classification 941
Personally Identifiable Information and Protected Health Information 943
PII 943
PHI 944
Privacy Enhancing Technologies 944
Roles and Responsibilities 945
Information Lifecycle 947
Impact Assessment 948
Terms of Agreement 948
Privacy Notice 949
Chapter Review Activities 949
Review Key Topics 949
Define Key Terms 949
Review Questions 950

Part VI: Final Preparation

Chapter 36 Final Preparation 953
Hands-on Activities 953
Suggested Plan for Final Review and Study 953
Summary 954

Glossary of Key Terms 955

Appendix A Answers to the "Do I Know This Already?" Quizzes and Review Questions 1023

Appendix B *CompTIA Security+ (SY0-601) Cert Guide* Exam Updates 1087

Index 1089

Online Elements:

Appendix C Study Planner

Glossary of Key Terms

About the Authors

Omar Santos is an active member of the cybersecurity community, where he leads several industry-wide initiatives. He is a best-selling author and trainer. Omar is the author of more than 20 books and video courses, as well as numerous white papers, articles, and security configuration guidelines and best practices. Omar is a principal engineer of the Cisco Product Security Incident Response Team (PSIRT), Security Research and Operations, where he mentors and leads engineers and incident managers during the investigation and resolution of cybersecurity vulnerabilities.

Omar co-leads the DEF CON Red Team Village, is the chair of the Common Security Advisory Framework (CSAF) technical committee, is the co-chair of the Forum of Incident Response and Security Teams (FIRST) Open Source Security working group, and has been the chair of several initiatives in the Industry Consortium for Advancement of Security on the Internet (ICASI). His active role helps businesses, academic institutions, state and local law enforcement agencies, and other participants dedicated to increasing the security of their critical infrastructures. You can find additional information about Omar's current projects at h4cker.org and can follow Omar on Twitter @santosomar.

Ron Taylor has been in the information security field for more than 20 years working in various areas focusing on both offense and defense security roles. Ten of those years were spent in consulting. In 2008, he joined the Cisco Global Certification Team as an SME in information assurance. From there, he moved into a position with the Security Research and Operations group, where his focus was mostly on penetration testing of Cisco products and services. He was also involved in developing and presenting security training to internal development and test teams globally, and provided consulting support to many product teams as an SME on product security testing. His next role was incident manager for the Cisco Product Security Incident Response Team (PSIRT). Currently, Ron is a security architect specializing in the Cisco security product line. He has held a number of industry certifications, including GPEN, GWEB, GCIA, GCIH, GWAPT, RHCE, CCSP, CCNA, CISSP, PenTest+, and MCSE. Ron has also authored books and video courses, teaches, and is involved in organizing a number of cybersecurity conferences, including the BSides Raleigh, Texas Cyber Summit, Grayhat, and the Red Team Village at DEFCON.

Twitter: @Gu5G0rman

Linkedin: www.linkedin.com/in/-RonTaylor

Joseph Mlodzianowski is an information security aficionado and adventurer; he started multiple villages at RSA Conference, DEFCON, and BLACK HAT, among others, including founding the Red Team Village with the help of great friends. He has been in the information technology security field for more than 25 years working in infrastructure, security, networks, systems, design, offense, and defense. Joseph is currently an enterprise security architect of Cisco Managed Services. He spent more than 10 years in the Department of Defense as an operator, principal security network engineer, and SME designing and deploying complex technologies and supporting missions around the world in multiple theaters. He has consulted, investigated, and provided support for multiple federal agencies over the past 15 years. Joseph continues to contribute to content, reviews, and editing in the certification testing and curriculum process. He spent almost 15 years in the energy sector supporting refineries, pipelines, and chemical plants; specializing in industrial control networks; and building data centers. Joseph holds a broad range of certifications, including the Cisco CCIE, CNE, CSNA, CNSS-4012, CISSP, ITILv4, NSA IAM, NSA IEM, OIAC1180, FEMA IS-00317, ACMA, First Responder, Hazmat Certified, Member of Bexar County Sheriff's Office CERT, MCSE, and Certified Hacking Investigator. He also is a founding contributor to the CyManII | Cybersecurity Manufacturing Innovation Institute, a member of Messaging Malware Mobile Anti-Abuse Working Group (M3aawg.org), and founder of the Texas Cyber Summit and Grayhat Conferences. He believes in giving back to the community and supporting nonprofits.

Twitter: @Cedoxx

Linkedin: www.linkedin.com/in/mlodzianowski/

Dedication

I would like to dedicate this book to my lovely wife, Jeannette, and my two beautiful children, Hannah and Derek, who have inspired and supported me throughout the development of this book.

—Omar

I would not be where I am today without the support of my family. Mom and Dad, you taught me the importance of work ethic and drive. Kathy, my wife of 20 years, you have supported me and encouraged me every step of the way. Kaitlyn, Alex, and Grace, you give me the strength and motivation to keep doing what I do.

—Ron

Without faith and spiritual guidance, none of us would be where we are. I would like to thank my Creator; Linda, my lovely wife of more than 20 years; and my daughter Lauren, for their unwavering support, patience, and encouragement while I work multiple initiatives and projects.

—Joseph

Acknowledgments

It takes a lot of amazing people to publish a book. Special thanks go to Chris Cleveland, Nancy Davis, Chris Crayton, and all the others at Pearson (and beyond) who helped make this book a reality. We appreciate everything you do!

About the Technical Reviewer

Chris Crayton is a technical consultant, trainer, author, and industry-leading technical editor. He has worked as a computer technology and networking instructor, information security director, network administrator, network engineer, and PC specialist. Chris has authored several print and online books on PC repair, CompTIA A+, CompTIA Security+, and Microsoft Windows. He has also served as technical editor and content contributor on numerous technical titles for several of the leading publishing companies. He holds numerous industry certifications, has been recognized with many professional and teaching awards, and has served as a state-level SkillsUSA final competition judge. Chris tech edited and contributed to this book to make it better for students and those wishing to better their lives.

We Want to Hear from You!

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

We welcome your comments. You can email or write to let us know what you did or didn't like about this book—as well as what we can do to make our books better.

Please note that we cannot help you with technical problems related to the topic of this book.

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

Email: community@informit.com

Reader Services

Register your copy of *CompTIA Security+ SY0-601 Cert Guide* at www.pearsonitcertification.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.pearsonit-certification.com/register and log in or create an account.* Enter the product ISBN 9780136770312 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Introduction

Welcome to the *CompTIA Security+ SY0-601 Cert Guide*. The CompTIA Security+ certification is widely accepted as the first security certification you should attempt to attain in your information technology (IT) career. The CompTIA Security+ certification is designed to be a vendor-neutral exam that measures your knowledge of industry-standard technologies and methodologies. It acts as a great stepping stone to other vendor-specific certifications and careers. We developed this book to be something you can study from for the exam and keep on your bookshelf for later use as a security resource.

We would like to note that it's unfeasible to cover all security concepts in depth in a single book. However, the Security+ exam objectives are looking for a basic level of computer, networking, and organizational security knowledge. Keep this in mind while reading through this text, and remember that the main goal of this text is to help you pass the Security+ exam, not to be the master of all security. Not just yet, at least!

Good luck as you prepare to take the CompTIA Security+ exam. As you read through this book, you will be building an impenetrable castle of knowledge, culminating in hands-on familiarity and the know-how to pass the exam.

Goals and Methods

The number one goal of this book is to help you pass the SY0-601 version of the CompTIA Security+ certification exam. To that effect, we have filled this book and practice exams with hundreds of questions/answers and explanations, including two full practice exams. The exams are located in Pearson Test Prep practice test software in a custom test environment. These tests are geared to check your knowledge and ready you for the real exam.

The CompTIA Security+ certification exam involves familiarity with computer security theory and hands-on know-how. To aid you in mastering and understanding the Security+ certification objectives, this book uses the following methods:

- **Opening topics list:** This list defines the topics to be covered in the chapter.
- **Foundation Topics:** The heart of the chapter. The text explains the topics from a theory-based standpoint, as well as from a hands-on perspective. This includes in-depth descriptions, tables, and figures that are geared to build your knowledge so that you can pass the exam. Each chapter covers a full objective from the CompTIA Security+ exam blueprint.
- **Key Topics:** The Key Topic icons indicate important figures, tables, and lists of information that you should know for the exam. They are interspersed throughout the chapter and are listed in table format at the end of the chapter.

- **Key Terms:** Key terms without definitions are listed at the end of each chapter. See whether you can define them, and then check your work against the complete key term definitions in the glossary.
- **Review Questions:** These quizzes and answers with explanation are meant to gauge your knowledge of the subjects. If an answer to a question doesn't come readily to you, be sure to review that portion of the chapter.
- **Practice Exams:** The practice exams are included in the Pearson Test Prep practice test software. These exams test your knowledge and skills in a realistic testing environment. Take them after you have read through the entire book. Master one; then move on to the next.

Who Should Read This Book?

This book is for anyone who wants to start or advance a career in computer security. Readers of this book can range from persons taking a Security+ course to individuals already in the field who want to keep their skills sharp or perhaps retain their job due to a company policy mandating they take the Security+ exam. Some information assurance professionals who work for the Department of Defense or have privileged access to DoD systems are required to become Security+ certified as per DoD directive 8570.1.

This book is also designed for people who plan on taking additional security-related certifications after the CompTIA Security+ exam. The book is designed in such a way to offer an easy transition to future certification studies.

Although not a prerequisite, it is recommended that CompTIA Security+ candidates have at least two years of IT administration experience with an emphasis on security. The CompTIA Network+ certification is also recommended as a prerequisite. Before you begin your Security+ studies, it is expected that you understand computer topics such as how to install operating systems and applications, and networking topics such as how to configure IP, what a VLAN is, and so on. The focus of this book is to show how to secure these technologies and protect against possible exploits and attacks. Generally, for people looking to enter the IT field, the CompTIA Security+ certification is attained after the A+ and Network+ certifications.

CompTIA Security+ Exam Topics

If you haven't downloaded the Security+ certification exam objectives, do it now from CompTIA's website: https://certification.comptia.org/. Save the PDF file and print it out as well. It's a big document; review it carefully. Use the exam objectives list and acronyms list to aid in your studies while you use this book.

The following tables are excerpts from the exam objectives document. Table I-1 lists the CompTIA Security+ domains and each domain's percentage of the exam.

Table I-1 CompTIA Security+ Exam Domains

Domain	Exam Topic	% of Exam
1.0	Attacks, Threats, and Vulnerabilities	24%
2.0	Architecture and Design	21%
3.0	Implementation	25%
4.0	Operations and Incident Response	16%
5.0	Governance, Risk, and Compliance	14%

The Security+ domains are then further broken down into individual objectives. Table I-2 lists the CompTIA Security+ exam objectives and their related chapters in this book. It does not list the bullets and sub-bullets for each objective.

Table I-2 CompTIA Security+ Exam Objectives

Objective	Chapter(s)
1.1 Compare and contrast different types of social engineering techniques.	1
1.2 Given a scenario, analyze potential indicators to determine the type of attack.	2
1.3 Given a scenario, analyze potential indicators associated with application attacks.	3
1.4 Given a scenario, analyze potential indicators associated with network attacks.	4
1.5 Explain different threat actors, vectors, and intelligence sources.	5
1.6 Explain the security concerns associated with various types of vulnerabilities.	6
1.7 Summarize the techniques used in security assessments.	7
1.8 Explain the techniques used in penetration testing.	8
2.1 Explain the importance of security concepts in an enterprise environment.	9
2.2 Summarize virtualization and cloud computing concepts.	10
2.3 Summarize secure application development, deployment, and automation concepts.	11
2.4 Summarize authentication and authorization design concepts.	12
2.5 Given a scenario, implement cybersecurity resilience.	13
2.6 Explain the security implications of embedded and specialized systems.	14
2.7 Explain the importance of physical security controls.	15
2.8 Summarize the basics of cryptographic concepts.	16
3.1 Given a scenario, implement secure protocols.	17

Objective	Chapter(s)
3.2 Given a scenario, implement host or application security solutions.	18
3.3 Given a scenario, implement secure network designs.	19
3.4 Given a scenario, install and configure wireless security settings.	20
3.5 Given a scenario, implement secure mobile solutions.	21
3.6 Given a scenario, apply cybersecurity solutions to the cloud.	22
3.7 Given a scenario, implement identity and account management controls.	23
3.8 Given a scenario, implement authentication and authorization solutions.	24
3.9 Given a scenario, implement public key infrastructure.	25
4.1 Given a scenario, use the appropriate tool to assess organizational security.	26
4.2 Summarize the importance of policies, processes, and procedures for incident response.	27
4.3 Given an incident, utilize appropriate data sources to support an investigation.	28
4.4 Given an incident, apply mitigation techniques or controls to secure an environment.	29
4.5 Explain the key aspects of digital forensics.	30
5.1 Compare and contrast various types of controls.	31
5.2 Explain the importance of applicable regulations, standards, or frameworks that impact organizational security posture.	32
5.3 Explain the importance of policies to organizational security.	33
5.4 Summarize risk management processes and concepts.	34
5.5 Explain privacy and sensitive data concepts in relation to security.	35

Companion Website

Register this book to get access to the Pearson Test Prep practice test software and other study materials plus additional bonus content. Check this site regularly for new and updated postings written by the authors that provide further insight into the more troublesome topics on the exam. Be sure to check the box that you would like to hear from us to receive updates and exclusive discounts on future editions of this product or related products.

To access this companion website, follow these steps:

1. Go to www.pearsonitcertification.com/register and log in or create a new account.

2. On your Account page, tap or click the **Registered Products** tab, and then tap or click the **Register Another Product** link.
3. Enter this book's ISBN (9780136770312).
4. Answer the challenge question as proof of book ownership.
5. Tap or click the **Access Bonus Content** link for this book to go to the page where your downloadable content is available.

Please note that many of our companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following the preceding steps, please visit http://www.pearsonitcertification.com/contact and select the "Site Problems/Comments" option. Our customer service representatives will assist you.

Pearson Test Prep Practice Test Software

As noted previously, this book comes complete with the Pearson Test Prep practice test software containing two full exams. These practice tests are available to you either online or as an offline Windows application. To access the practice exams that were developed with this book, please see the instructions in the card inserted in the sleeve in the back of the book. This card includes a unique access code that enables you to activate your exams in the Pearson Test Prep software.

NOTE The cardboard sleeve in the back of this book includes a piece of paper. The paper lists the activation code for the practice exams associated with this book. Do not lose the activation code. On the opposite side of the paper from the activation code is a unique, one-time-use coupon code for the purchase of the Premium Edition eBook and Practice Test.

Accessing the Pearson Test Prep Software Online

The online version of this software can be used on any device with a browser and connectivity to the Internet including desktop machines, tablets, and smartphones. To start using your practice exams online, simply follow these steps:

1. Go to www.PearsonTestPrep.com and select **Pearson IT Certification** as your product group.
2. Enter your email/password for your account. If you do not have an account on PearsonITCertification.com or InformIT.com, you will need to establish one by going to PearsonITCertification.com/join.

3. On the My Products tab, tap or click the **Activate New Product** button.
4. Enter this book's activation code and click **Activate**.
5. The product will now be listed on your My Products tab. Tap or click the **Exams** button to launch the exam settings screen and start your exam.

Accessing the Pearson Test Prep Software Offline

If you wish to study offline, you can download and install the Windows version of the Pearson Test Prep software. There is a download link for this software on the book's companion website, or you can just enter this link in your browser:

http://www.pearsonitcertification.com/content/downloads/pcpt/engine.zip

To access the book's companion website and the software, simply follow these steps:

1. Register your book by going to http://www.pearsonitcertification.com/register and entering the ISBN: 9780136770312.
2. Respond to the challenge questions.
3. Go to your account page and select the **Registered Products** tab.
4. Click the **Access Bonus Content** link under the product listing.
5. Click the **Install Pearson Test Prep Desktop Version** link under the Practice Exams section of the page to download the software.
6. Once the software finishes downloading, unzip all the files on your computer.
7. Double-click the application file to start the installation, and follow the onscreen instructions to complete the registration.
8. Once the installation is complete, launch the application and click the **Activate Exam** button on the My Products tab.
9. Click the **Activate a Product** button in the Activate Product Wizard.
10. Enter the unique access code found on the card in the sleeve in the back of your book and click the **Activate** button.
11. Click **Next** and then the **Finish** button to download the exam data to your application.
12. You can now start using the practice exams by selecting the product and clicking the **Open Exam** button to open the exam settings screen.

Note that the offline and online versions will synch together, so saved exams and grade results recorded on one version will be available to you on the other as well.

Customizing Your Exams

Once you are in the exam settings screen, you can choose to take exams in one of three modes:

- Study Mode
- Practice Exam Mode
- Flash Card Mode

Study Mode enables you to fully customize your exams and review answers as you are taking the exam. This is typically the mode you would use first to assess your knowledge and identify information gaps. Practice Exam Mode locks certain customization options, as it is presenting a realistic exam experience. Use this mode when you are preparing to test your exam readiness. Flash Card Mode strips out the answers and presents you with only the question stem. This mode is great for late-stage preparation when you really want to challenge yourself to provide answers without the benefit of seeing multiple-choice options. This mode will not provide the detailed score reports that the other two modes will, so it should not be used if you are trying to identify knowledge gaps.

In addition to these three modes, you will be able to select the source of your questions. You can choose to take exams that cover all of the chapters or you can narrow your selection to just a single chapter or the chapters that make up specific parts in the book. All chapters are selected by default. If you want to narrow your focus to individual chapters, simply deselect all the chapters and then select only those on which you wish to focus in the Objectives area.

You can also select the exam banks on which to focus. Each exam bank comes complete with a full exam of questions that cover topics in every chapter. You can have the test engine serve up exams from all banks or just from one individual bank by selecting the desired banks in the exam bank area.

There are several other customizations you can make to your exam from the exam settings screen, such as the time of the exam, the number of questions served up, whether to randomize questions and answers, whether to show the number of correct answers for multiple-answer questions, or whether to serve up only specific types of questions. You can also create custom test banks by selecting only questions that you have marked or questions on which you have added notes.

Updating Your Exams

If you are using the online version of the Pearson Test Prep software, you should always have access to the latest version of the software as well as the exam data. If you are using the Windows desktop version, every time you launch the software, it will check to see if there are any updates to your exam data and automatically download

any changes that were made since the last time you used the software. This requires that you are connected to the Internet at the time you launch the software.

Sometimes, due to many factors, the exam data may not fully download when you activate your exam. If you find that figures or exhibits are missing, you may need to manually update your exams.

To update a particular exam you have already activated and downloaded, simply select the **Tools** tab and click the **Update Products** button. Again, this is only an issue with the desktop Windows application.

If you wish to check for updates to the Pearson Test Prep exam engine software, Windows desktop version, simply select the **Tools** tab and click the **Update Application** button. This will ensure you are running the latest version of the software engine.

Premium Edition eBook and Practice Tests

This book also includes an exclusive offer for 80 percent off the Premium Edition eBook and Practice Tests edition of this title. Please see the coupon code included with the cardboard sleeve for information on how to purchase the Premium Edition.

Figure Credits

Cover image: TippaPatt/Shutterstock
Chapter opener image: Charlie Edwards/Photodisc/Getty Images
Figures 4-2 and 4-3 courtesy of Cisco Systems, Inc
Figure 5-1 © 2015-2021, The MITRE Corporation
Figure 7-5 courtesy of Cisco Systems, Inc
Figures 10-1, 10-7, 10-10 courtesy of Cisco Systems, Inc
Figure 10-8 © 2021, Amazon Web Services, Inc
Figure 12-3 courtesy of Secret Double Octopus
Figure 12-4 courtesy of Active-Directory-FAQ
Figure 12-5 courtesy of Robert Koczera/123RF
Figures 13-1 and 13-2 © AsusTek Computer Inc.
Figure 14-1 Raspberry Pi courtesy of handmadepictures/123RF
Figure 14-3 courtesy of CSS Electronics
Figure 14-4 courtesy of strajinsky/Shutterstock
Figure 14-5 courtesy of RingCentral
Figures 14-6 and 15-4 from rewelda/Shutterstock
Figure 15-1 courtesy of Kyryl Gorlov/123RF
Figure 15-2 courtesy of Aliaksandr Karankevich/123RF
Figures 16-9 and 16-10 courtesy of ssl2buy.com
Figure 17-1 courtesy of hostinger.com
Figure 17-3 courtesy of wiki.innovaphone.com
Figure 17-4 courtesy of Adaptive Digital Technologies
Figure 18-4 © Microsoft 2021
Figure 18-5 courtesy of Microsoft Corporation
Figure 18-7 courtesy of Checkmarx Ltd
Figure 19-1 courtesy of Cisco Systems, Inc
Figures 19-5, 19-8 through 19-11, 19-15, and 19-16 courtesy of Cisco Systems, Inc
Figure 19-21 © Microsoft 2021
Figures 20-2 and Figure 24-4 © Microsoft 2021
Figure 20-4 © D-Link Corporation
Figure 20-5 courtesy of Cisco Systems, Inc
Figures 21-1 and 21-2 © 1992-2020 Cisco
Figures 23-1 and 23-2 courtesy of Cisco Systems, Inc
Figures 23-5 through 23-9 © Microsoft 2021
Figures 24-2 and 24-5 courtesy of Cisco Systems, Inc
Figures 24-6 through 24-10 © Microsoft 2021
Figures 25-1 and 25-2 ©1998–2021 Mozilla Foundation
Figure 26-1 © OffSec Services Limited 2021
Figures 26-2, 26-5, 26-6, 26-10 through 26-15, 26-18, 26-19 © 2021 The Linux Foundation
Figures 26-3, 26-4, 26-7 through 26-9 © Microsoft 2021
Figure 26-16 © 2021 Tenable, Inc
Figure 26-17 © 2010-2020, Cuckoo Foundation
Figure 26-21 © Wireshark Foundation
Figure 26-22 © X-Ways Software Technology AG
Figure 27-3 courtesy of Cisco Systems, Inc
Figure 27-4 courtesy of Evolve IP, LLC
Figure 28-1 © 2021 Tenable, Inc

Figure 28-2 © 2021 LogRhythm, Inc
Figure 28-3 © MaxBelkov
Figure 28-4 © 1992-2020 Cisco
Figure 28-5, 28-6, 28-10, 28-11, and 28-18 © Microsoft 2021
Figures 28-7 through 28-9, 28-14, and 28-15 © 2021 The Linux Foundation
Figures 28-12 and 28-13 © 1992-2020 Cisco
Figure 28-16 © 2021 NXLog Ltd
Figure 28-20 © 2021 SolarWinds Worldwide, LLC
Figure 28-21 © 2003-2021 sFlow.org
Figure 28-22 © Plixer, LLC
Figure 28-23 © Microsoft 2021
Figures 29-1 and 29-2 © Microsoft 2021
Figure 33-1 © Microsoft 2021



This chapter covers the following topics related to Objective 1.1 (Compare and contrast different types of social engineering techniques) of the CompTIA Security+ SY0-601 certification exam:

- Phishing
- Smishing
- Vishing
- Spam
- Spam over internet messaging (SPIM)
- Spear phishing
- Dumpster diving
- Shoulder surfing
- Pharming
- Tailgating
- Eliciting information
- Whaling
- Prepending
- Identity fraud
- Invoice scams
- Credential harvesting
- Reconnaissance
- Hoax
- Impersonation
- Watering hole attack
- Typo squatting
- Pretexting
- Influence campaigns (hybrid warfare and social media)
- Principles of social engineering and reasons for effectiveness (including Authority, Intimidation, Consensus, Scarcity, Familiarity, Trust, and Urgency)

Comparing and Contrasting Different Types of Social Engineering Techniques

Social engineering attacks leverage the weakest link, which is typically the human user. If an attacker can get a user to reveal information, it is much easier for the attacker to cause harm than it is by using some other method of reconnaissance. Social engineering can be accomplished through email or misdirection of web pages, prompting a user to click something that leads to the attacker gaining information. Social engineering can also be done in person by an insider or an outside entity or over the phone.

This chapter delves into the methods and techniques that social engineers can employ to gain access to buildings and systems and obtain company data and personal information. It also covers the various ways that these social engineers can be defeated.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 1-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 1-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Social Engineering Fundamentals	1–9
User Security Awareness Education	10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. An attacker sends a targeted email with a malicious attachment to a user in your company. This attacker researched public information about the user to send a "more personal" and targeted email to the user. Which of the following is this type of attack?
   a. Spear phishing
   b. Typo squatting
   c. Pharming
   d. None of these answers are correct.

2. Which of the following is an example of a tool that can be used specifically to perform social engineering attacks?
   a. Maltego
   b. SET
   c. The Harvester
   d. Recon-NG

3. Which of the following best describes the difference between smishing and vishing?
   a. Vishing is a social engineering attack in which the attacker calls the user over the phone and then persuades the user to reveal sensitive information or perform a given action. Smishing is a type of phishing campaign using SMS text messages instead of email.
   b. Vishing is a social engineering attack in which the attacker leaves a voicemail and then persuades the user to reveal sensitive information or perform a given action. Smishing is a type of typo squatting and pharming campaign using Bluetooth.
   c. Vishing is a social engineering attack in which the attacker leaves a voicemail and then persuades the user to reveal sensitive information or perform a given action. Smishing is a type of typo squatting and pharming campaign using short Internet messaging systems.
   d. None of these answers are correct.

4. A(n) _____ is a small space that can usually only fit one person, used to combat tailgating.
   a. tunnel-gap
   b. tunnel-trap
   c. piggyback
   d. access control vestibule

5. Which of the following best describes pretexting?
   a. Impersonation
   b. Social engineering
   c. Whaling
   d. Pharming

6. Which of the following refers to the act of incorporating malicious ads on trusted websites, which results in users' browsers being inadvertently redirected to sites hosting malware?
   a. Malvertising
   b. Pharming
   c. Active ad exploitation
   d. Whaling

7. Which of the following is true about spear phishing?
   a. Spear phishing attacks use the Windows Administrative Center.
   b. Spear phishing attacks are phishing attempts that are constructed in a very specific way and directly targeted to specific individuals or companies.
   c. Spear phishing, whaling, and phishing are the same type of attack.
   d. Spear phishing attacks use the Windows PowerShell.

8. Derek is the CEO of a Fortune 500 company. He received an email with a malicious attachment. Once Derek clicked on the attachment, malware was installed on his system. Which of the following best describes this attack?
   a. Smishing
   b. Vishing
   c. Whaling
   d. Pretexting

9. Which of the following is true about social engineering motivation techniques?
    a. Social proof can be used to create a feeling of urgency in a decision-making context. It is possible to use specific language in an interaction to present a sense of urgency and manipulate the victim.
    b. Scarcity can be used to create a feeling of urgency in a decision-making context. It is possible to use specific language in an interaction to present a sense of urgency and manipulate the victim.
    c. Scarcity cannot be used to create a feeling of urgency in a decision-making context. It is possible to use specific language in an interaction to present a sense of urgency and manipulate your victim.
    d. Social proof cannot be used in an interrogation because it is illegal. It is not legal to use specific language in an interaction to present a sense of urgency and manipulate your victim.

10. Which of the following recommendations can be used in user security awareness education?
    a. Adhere to the organization's clean desk policy, which states that all documents, electronics, personally owned devices, and other items be put away (or locked away) when the user is not at his or her desk, or other work area.
    b. Always screen your email and phone calls carefully and keep a log of events. This is also known as communications vetting.
    c. Use encryption when possible to protect emails, phone calls, and data.
    d. All of these answers are correct.

Foundation Topics

Social Engineering Fundamentals

Let's discuss a low note in our society. Because that is what information security–based social engineering is—a low form of behavior, but an effective one. It is estimated that 1 out of 10 people is conned every year through social engineering methods, and as many as half of them don't even know it has occurred. It's glorified in the movies, but in real life it can have devastating consequences to an organization and to innocent individuals.

Social engineering is the act of manipulating users into revealing confidential information or performing other actions detrimental to the user. Examples of social engineering are common in everyday life. A basic example would be a person asking for your username and password over the phone; often the person uses flattery to gain information. Malicious people use various forms of social engineering in an attempt to steal whatever you have of value: your money, information, identity, confidential company data, or IT equipment.

A primary example is attackers leveraging normal user behavior. Suppose that you are a security professional who is in charge of the network firewalls and other security infrastructure equipment in your company. An attacker could post a job offer for a very lucrative position and make it very attractive to you, the victim. Suppose the job description lists benefits and compensation far beyond what you are already making at your company. You decide to apply for the position. The criminal (attacker) then schedules an interview with you. Because you are likely to "show off" your skills and work, the attacker may be able to get you to explain how you have configured the firewalls and other network infrastructure devices for your company. You might disclose information about the firewalls used in your network, how you have configured them, how they were designed, and so on. Your answers give the attacker a lot of knowledge about the organization without requiring the attacker to perform any type of scanning or reconnaissance on the network.

Let's take a look at another example: suppose that you are a security guard and a pregnant woman comes to you saying that she is feeling sick and that she needs a bathroom immediately. You are courteous and escort her to the bathroom inside your premises, where she then uses her laptop to connect to your Wi-Fi or put a wireless rogue access point (AP) to lure some of your users to connect to such an AP. We are good and kind people (most of the time) and will continue to be the weakest link in the cybersecurity world because this human nature becomes our weakness.

Social engineers will rely on information. For example, open-source intelligence (OSINT) is a way that attackers can gain knowledge about a target. OSINT includes

media, public government data, commercial data, and academic publications. You will learn additional details about OSINT in Chapter 5, "Understanding Different Threat Actors, Vectors, and Intelligence Sources."

TIP You can find several references to tools that can be used for OSINT reconnaissance at my GitHub repository: https://github.com/The-Art-of-Hacking/h4cker. This GitHub repository has thousands of cybersecurity references that can be very helpful for this certification and many others.

The main reason that social engineering succeeds is a lack of user awareness. But social engineering can also be effective in environments in which the IT personnel have little training, and in public areas—for example, public buildings with shared office space. Let's discuss some of the more common types of social engineering.

Common social engineering techniques include the following:

- Phishing and spear phishing
- Smishing
- Vishing
- Spam
- Spam over Internet messaging (SPIM)
- Dumpster diving
- Shoulder surfing
- Pharming
- Tailgating
- Eliciting information
- Whaling
- Prepending
- Identity fraud
- Invoice scams
- Credential harvesting
- Reconnaissance
- Hoax
- Impersonation or pretexting

- Eavesdropping
- Baiting
- Watering hole attack
- Typo squatting

Key Topic

Phishing and Spear Phishing

Phishing is the attempt at fraudulently obtaining private information. A phisher usually masquerades as someone else, perhaps another entity. There are two main differences between phishing and pretexting. ***Pretexting*** is the act of impersonating or "spoofing" someone else's identity. For example, an attacker could impersonate an employee of a company or a business partner to attempt to steal sensitive data from the victim. Phishing is usually done by electronic communication, not in person. Second, little information about the target is necessary. A phisher might target thousands of individuals without much concern as to their background. An example of phishing would be an email that requests verification of private information. The email probably leads to a malicious website designed to lure people into a false sense of security to fraudulently obtain information. The website often looks like a legitimate website. A common phishing technique is to pose as a vendor (such as an online retailer or domain registrar) and send the target email confirmations of orders that they supposedly placed.

This is a triple-whammy. First, the orders are obviously fake; a person might say, "Hey, wait! I didn't place these orders!" and perhaps click the link(s) in the email, leading the person to the false web page. Second, if a person thinks it's a legitimate order (perhaps the person does many orders, and the fraudulent one looks like another legitimate one), the person might click a link to track the order, again leading to the bogus web page. Third, once at the web page, the person is asked to enter credentials for an account (which then leads to credit card fraud and ID theft), and in addition to that the page might have Trojans and other malicious scripts that are delivered to the unsuspecting person on exit. Sheesh, talk about cyberbullying!

Generally, no information about the target is necessary for a phishing attack. However, some "phishermen" actually target specific groups of people or even specific individuals. This is known as ***spear phishing***. And when an attacker targets key stakeholders and senior executives (CEOs, CFOs, and so on), it is known as ***whaling***. Whaling attacks are much more detailed and require that the attacker know a good deal of information about the target (much of which is freely available on the Internet).

Many different types of social engineering are often lumped into what is referred to as phishing, but actual phishing for private information is normally limited to email and websites.

Several open-source tools can help automate or accelerate a social engineering attack (including sending phishing and spear phishing emails). An example of these tools (and one of the most popular) is the Social Engineering Toolkit (SET). Example 1-1 demonstrates how attackers can use SET to send spear phishing emails to a victim and either select existing payloads or create their own.

Example 1-1 The Social Engineering Toolkit

```
[---]         The Social-Engineer Toolkit (SET)          [---]
[---]         Created by: David Kennedy (ReL1K)          [---]
                        Version: 8.0.3
                     Codename: 'Maverick'
[---]         Follow us on Twitter: @TrustedSec          [---]
[---]         Follow me on Twitter: @HackingDave         [---]
[---]        Homepage: https://www.trustedsec.com        [---]
         Welcome to the Social-Engineer Toolkit (SET).
          The one stop shop for all of your SE needs.
   The Social-Engineer Toolkit is a product of TrustedSec.
            Visit: https://www.trustedsec.com
   It's easy to update using the PenTesters Framework! (PTF)
Visit https://github.com/trustedsec/ptf to update all your tools!

 Select from the menu:
   1) Spear-Phishing Attack Vectors
   2) Website Attack Vectors
   3) Infectious Media Generator
   4) Create a Payload and Listener
   5) Mass Mailer Attack
   6) Arduino-Based Attack Vector
   7) Wireless Access Point Attack Vector
   8) QRCode Generator Attack Vector
   9) Powershell Attack Vectors
  10) Third Party Modules
  99) Return back to the main menu.
set> 1

The Spearphishing module allows you to specially craft email messages
and send  them to a large (or small) number of people with attached
fileformat malicious  payloads. If you want to spoof your email
address, be sure "Sendmail" is installed (apt-get install sendmail)
and change the config/set_config SENDMAIL=OFF  flag to SENDMAIL=ON.
```

```
There are two options, one is getting your feet wet and letting SET
do everything for you (option 1), the second is to create your own
FileFormat payload and use it in your own attack. Either way, good
luck and enjoy!

   1) Perform a Mass Email Attack
   2) Create a FileFormat Payload
   3) Create a Social-Engineering Template
  99) Return to Main Menu
set:phishing>1
/usr/share/metasploit-framework/
 Select the file format exploit you want.
 The default is the PDF embedded EXE.
           ********** PAYLOADS **********
   1) SET Custom Written DLL Hijacking Attack Vector (RAR, ZIP)
   2) SET Custom Written Document UNC LM SMB Capture Attack
   3) MS15-100 Microsoft Windows Media Center MCL Vulnerability
   4) MS14-017 Microsoft Word RTF Object Confusion (2014-04-01)
   5) Microsoft Windows CreateSizedDIBSECTION Stack Buffer Overflow
   6) Microsoft Word RTF pFragments Stack Buffer Overflow (MS10-087)
   7) Adobe Flash Player "Button" Remote Code Execution
   8) Adobe CoolType SING Table "uniqueName" Overflow
   9) Adobe Flash Player "newfunction" Invalid Pointer Use
  10) Adobe Collab.collectEmailInfo Buffer Overflow
  11) Adobe Collab.getIcon Buffer Overflow
  12) Adobe JBIG2Decode Memory Corruption Exploit
  13) Adobe PDF Embedded EXE Social Engineering
  14) Adobe util.printf() Buffer Overflow
  15) Custom EXE to VBA (sent via RAR) (RAR required)
  16) Adobe U3D CLODProgressiveMeshDeclaration Array Overrun
  17) Adobe PDF Embedded EXE Social Engineering (NOJS)
  18) Foxit PDF Reader v4.1.1 Title Stack Buffer Overflow
  19) Apple QuickTime PICT PnSize Buffer Overflow
  20) Nuance PDF Reader v6.0 Launch Stack Buffer Overflow
  21) Adobe Reader u3D Memory Corruption Vulnerability
  22) MSCOMCTL ActiveX Buffer Overflow (ms12-027)
set:payloads>
```

NOTE SET and other penetration testing tools are not covered in the exam. The instance provided in Example 1-1 is for your reference only. SET and similar tools can be installed in different Linux distributions. Some of the most popular Linux distributions for security penetration testing are Kali Linux, Parrot Security, and Black Arch.

To defend against this, users should install a phishing filter or add-on and enable it on the web browser. Also, a person should be trained to realize that institutions will *not* call or email requesting private information. If people are not sure, they should hang up the phone or simply delete the email. A quick way to find out whether an email is phishing for information is to hover over a link. You will see a URL domain name that is far different from that of the institution that the phisher is claiming to be, probably a URL located in a distant country. Many of these phishers are also probably engaging in spy-phishing: a combination of spyware and phishing that effectively makes use of spyware applications. A spyware application of this sort is downloaded to the target, which then enables additional phishing attempts that go beyond the initial phishing website.

Key Topic

Smishing

Because phishing has been an effective tactic for threat actors, they have found ways other than using email to fool their victims into following malicious links or activating malware from emails. A number of phishing campaigns have used Short Message Service (SMS) to send malware or malicious links to mobile devices. This social engineering technique is also known as ***smishing*** (short for SMS phishing).

One example of SMS phishing is the Bitcoin-related SMS scams that have surfaced in recent years. Numerous victims have received messages instructing them to click links to confirm their accounts and claim Bitcoins. When a user clicks such a link, he or she might be fooled into entering sensitive information on that attacker's site.

Key Topic

Vishing

Phishing attacks can also be accomplished by telephone conversations. Phone or voice phishing, known as ***vishing***, works in the same manner as phishing but is initiated by a phone call (often using VoIP systems). Vishing is typically used to steal sensitive information such as Social Security numbers, credit card numbers, or other information used in identity theft schemes. Attackers might impersonate and spoof caller ID to obfuscate themselves when performing voice phishing attacks.

In many vishing attacks, the phone call often sounds like a prerecorded message from a legitimate institution (bank, online retailer, donation collector, and so on).

The message asks the unsuspecting person for confidential information such as name, bank account numbers, codes, and so on—all under the guise of needing to verify information for the person's protection. It's really the opposite, of course, and many people are caught unaware by these types of scams every day. Through the use of automated systems (such as the ones telemarketers use), vishing can be perpetuated on large groups of people with little effort.

NOTE A similar technique using automated systems is known as *war-dialing*. This type of attack occurs when a device (modem or other system) is used to scan a list of telephone numbers and dial them in search of computer systems and fax machines. The technique sifts out the phone numbers associated with voice lines and the numbers associated with computers. The result is a list that can later be used by other attackers for various purposes.

Key Topic

Spam and Spam over Internet Messaging (SPIM)

Spam has traditionally been the term used to describe unwanted and unsolicited email messages. However, email messages are not the only way that attackers or malicious users send spam. Users may also receive spam (numerous messages) via social media sites and instant/Internet messaging applications. For instance, you can be talking to other security professionals over a Discord server and someone may hop on the server and send dozens of unwanted messages, sometimes protesting against something. A better example is receiving unwanted messages and scam offers via your Facebook Messenger, WhatsApp, Keybase, Signal, WeChat, Telegram, or any other instant messaging application.

TIP Several free services enable you to obtain DNS-based blocking lists of spammers, as well as ways that you can report spam events. An example is SpamCop (https://www.spamcop.net).

Key Topic

Dumpster Diving

When a person literally scavenges for private information in garbage and recycling containers, that is known as ***dumpster diving***. Any sensitive documents should be stored in a safe place as long as possible. When they are no longer necessary, they should be shredded. (Some organizations incinerate their documents.) Information might be found not only on paper but also on hard drives or removable media. Chapter 15, "Understanding the Importance of Physical Security Controls," covers proper recycling and/or destruction of hard drives in more detail.

Key Topic

Shoulder Surfing

Shoulder surfing occurs when a person uses direct observation to find out a target's password, PIN, or other such authentication information. The simple resolution for this type of behavior is for the user to shield the screen with privacy screens or filters, keypads, or other authentication-requesting devices. A more aggressive approach is to courteously ask the suspected shoulder surfer to move along. Also, private information should never be left on a desk or out in the open. Computers should be locked or logged off when the user is not in the immediate area. From a more technical perspective, password masking can be implemented (if not already), where typed passwords appear only as asterisks or dots on the screen. Always check if your systems, applications, and devices use password masking. Some lesser devices (such as SOHO routers) may not implement password masking by default, and that might go against company policy due to the inherent lack of security. Shoulder surfing, along with eavesdropping, and dumpster diving are examples of no-tech hacking.

Key Topic

Pharming

Pharming is the term used to describe a threat actor redirecting a victim from a valid website or resource to a malicious one that could be made to appear as the valid site to the user. From there, an attempt is made to extract confidential information from the user or to install malware in the victim's system. Pharming can be done by altering the hosts file on a victim's system, through DNS poisoning, or by exploiting a vulnerability in a DNS server.

Figure 1-1 illustrates the mechanics of how pharming works.

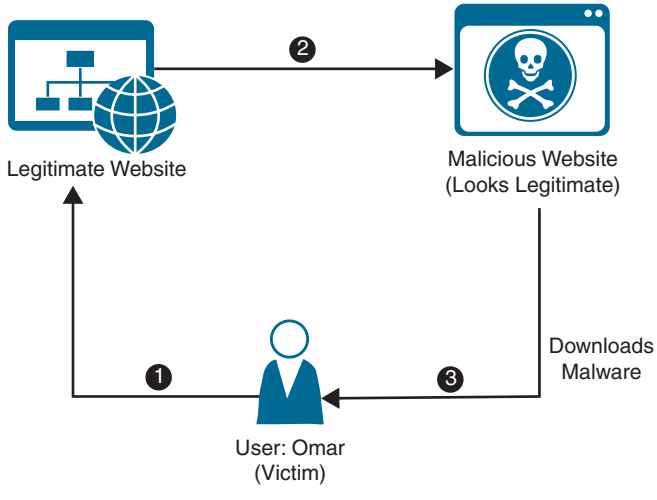


FIGURE 1-1 Pharming Example

The following steps are illustrated in Figure 1-1:

1. The user (Omar) visits a legitimate website and clicks a legitimate link.
2. Omar's system is compromised, the hosts file is modified, and Omar is redirected to a malicious site that appears to be legitimate. (This could also be accomplished by compromising a DNS server or spoofing a DNS reply.)
3. Malware is downloaded and installed on Omar's system.

Key Topic

Piggybacking or Tailgating

When an unauthorized person tags along with an authorized person to gain entry to a restricted area—usually with the person's consent—that person is ***piggybacking***. ***Tailgating*** is essentially the same with one difference: it is usually without the authorized person's consent. Both of these attacks can be defeated through the use of access control vestibules (formerly known as mantraps). An access control vestibule is a small space that can usually fit only one person. It has two sets of interlocking doors; the first set must be closed before the other will open, creating a sort of waiting room where people are identified (and cannot escape!). This technique is often used in server rooms and data centers. Multifactor authentication is often used in conjunction with an access control vestibule; for example, using a proximity card and PIN at the first door and using a biometric scan at the second. An access control vestibule is an example of a preventive security control. Turnstiles, double entry doors, and security guards are other less expensive (and less effective) solutions to the problem of piggybacking and tailgating and help address confidentiality in general.

Key Topic

Eliciting Information

Key components of social engineering are how someone influences, interrogates, and impersonates others. In short, ***eliciting information*** is the act of gaining knowledge or information from people. In most cases, an attacker gets information from the victim without directly asking for that particular information.

How an attacker interrogates and interacts with a victim is crucial for the success of the social engineering campaign. An interrogator can ask good open-ended questions to learn about an individual's viewpoints, values, and goals. The interrogator can then use any information revealed to continue to gather additional information or to obtain information from another victim.

It is also possible for an interrogator to use closed-ended questions to get more control of the conversation and to lead the conversation or to stop the conversation. Asking too many questions can cause the victim to shut down the interaction, and asking too few questions might seem awkward. Successful social engineering

interrogators use a narrowing approach in their questioning to gain the most information from the victim.

Interrogators pay close attention to the following:

- The victim's posture or body language
- The coloring of the victim's skin, such as the victim's face color becoming pale or red
- The direction of the victim's head and eyes
- Movement of the victim's hands and feet
- The victim's mouth and lip expressions
- The pitch and rate of the victim's voice, as well as changes in the voice
- The victim's words, including their length, the number of syllables, dysfunctions, and pauses

With pretexting—or impersonation—an attacker presents as someone else in order to gain access to information. In some cases, this type of social engineering attack can be very simple, such as quickly pretending to be someone else within an organization; in other cases, it can involve creating a whole new identity and then using that identity to manipulate the receipt of information. Social engineers might use pretexting to impersonate individuals in certain jobs and roles, even if they do not have experience in those jobs or roles.

For example, a social engineer might impersonate an IT support worker and provide unsolicited help to a user. Impersonating IT staff can be very effective because if you ask someone if he or she has a technical problem, it is quite likely that the victim will think about it and say something like, "Yes, as a matter of fact…yesterday, this weird thing happened to my computer." Impersonating IT staff can give an attacker physical access to systems in the organization. The attacker who has physical access can use a USB stick containing custom scripts to compromise a computer within seconds.

Key Topic

Whaling

Whaling is similar to phishing and spear phishing; however, with whaling, the attack is targeted at high-profile business executives and key individuals in a corporation. So, what is the difference between whaling and spear phishing? Like threat actors conducting spear phishing attacks, threat actors conducting whaling attacks also create emails and web pages to serve malware or collect sensitive information; however, the whaling attackers' emails and pages have a more official or serious look and feel. Whaling emails are designed to look like critical business emails or something from

someone who has legitimate authority, either externally or even internally in the company itself. In whaling attacks, web pages are designed to specifically address high-profile victims. In a regular phishing attack, the email might be a faked warning from a bank or service provider. In whaling attacks, the email or a web page would be created with a more serious executive-level form. The content is created to target an upper manager, such as the CEO, or an individual who might have credentials for valuable accounts within the organization. In summary, a whaling attack takes additional steps to target and entice higher-profile victims.

The main goal in whaling attacks is to steal sensitive information or compromise the victim's system and then target other key high-profile victims.

Attackers could use multifaceted attacks (also known as combined social engineering attacks). For instance, an attacker could send a spear-phishing email to a victim and then follow up with a phone call. This makes the attack even more effective.



Prepending

You can often configure your email servers or email cloud services to prepend a message in the email subject line to identify emails that are coming from outside of the organization. This technique is also known as ***prepending***. For example, the subject line of an email could read something like this:

```
Subject: "[EXTERNAL] Order Update"
```



Identity Fraud

Social engineering has been used by many threat actors and scammers to steal sensitive information to then open new bank accounts, make purchases, and get tax refunds, among many other scenarios. The United States Federal Trade Commission (FTC) created an educational website to help you report and recover from identity theft at https://www.identitytheft.gov. Other countries have similar websites and resources.

As stated by the United States FTC, "There are many ways that you might discover that someone is using your information. You might get a notice from the IRS or find unfamiliar accounts on your credit report. You might notice strange withdrawals from your bank account, get bills that aren't yours, or get calls about debts that you don't owe."



Invoice Scams

Many scammers continue to send unsolicited emails and other messages to victims that include "an invoice" (as a malicious attachment) for something that they

have not purchased. For instance, a scammer might send you a targeted email message including an invoice for something that you most likely have purchased in the past. The scammer may look for information about you (using OSINT) and then send a targeted email with such an invoice. For example, the attacker/scammer may have found your LinkedIn profile and then will send an email with a fake invoice from LinkedIn. Because you use that service, you might be more susceptible to click on that attachment and, subsequently, for your system to get compromised.

Key Topic

Credential Harvesting

Nowadays, you hear about a new breach pretty much on a daily basis. In many cases, attackers dump user credentials in sites like Pastebin, GitHub, and many others. Consequently, it is easier for attackers to compromise systems because users reuse their usernames and passwords to authenticate to other sites and systems. ***Credential harvesting*** (often called account harvesting) refers to the attacking technique or activities of grabbing legitimate usernames and even passwords to gain access to systems to steal information or to use them for malicious purposes.

Key Topic

Reconnaissance

Threat actors obtain many details about their victims before launching an attack. They can do this by actively performing scans of systems and networks (active reconnaissance). However, threat actors can also leverage open-source intelligence (OSINT) or human intelligence (HUMINT) to potentially obtain sensitive information from a target without the need of using active scans or sending any packets to the targeted network and underlying systems (this is referred to as passive reconnaissance).

TIP The MITRE Corporation (MITRE) is a United States government-funded research organization. It has created different standards and frameworks that are used by many organizations and cybersecurity professionals. MITRE created a set of matrices (called the MITRE ATT&CK) that document the tactics and techniques attackers use to compromise networks and systems. As part of the same effort, MITRE maintains a list of tactics and techniques used by attackers to gather information about their victims before compromising their systems and networks (PRE-ATT&CK). MITRE's ATT&CK can be accessed at https://attack.mitre.org and PRE-ATT&CK can be accessed at https://attack.mitre.org/matrices/enterprise/pre/.

Hoaxes

A ***hoax*** is the attempt at deceiving people into believing something that is false. The differences between hoaxes and phishing can be quite gray. However, hoaxes can come in person or through other means of communication, whereas phishing is generally relegated to e-communication and phone. Although phishing can occur at any time, and with the specific goal of obtaining private information, a hoax can often be perpetuated on holidays or other special days and could be carried out simply for fun. Regardless, hoaxes can use up valuable organization resources: email replies, Internet bandwidth used, time spent, and so on. An example of a "harmless" hoax was Google's supposed name change to "Topeka" on April Fools' Day 2010. An example of a financially harmful hoax was the supposed assassination of Bill Gates on April Fools' Day 2003. This hoax led to stock market fluctuations and loss of profit in Asia. Some companies place a time limit on jokes and hoaxes indicating that the affected person has become nonproductive; for example, 3 percent of the workday.

Pretexting, malicious insider attempts, diversion theft, phishing, and hoaxes are all known as *confidence tricks*, thus the term *con*, and are committed by "bunko" artists. However, there are even lower ways to get access to people's information; these often are used with the previous methods. They include shoulder surfing, eavesdropping, dumpster diving, baiting, and piggybacking.

Impersonation or Pretexting

Pretexting is a type of social engineering attack in which a person invents a scenario, or pretext, in the hope of persuading a victim to divulge information. Preparation and some prior information are often needed before attempting a pretext; impersonation is often a key element. By impersonating the appropriate personnel or third-party entities, a person performing a pretext hopes to obtain records about an organization, its data, and its personnel. IT people and employees should always be on the lookout for impersonators and always ask for identification. If there is any doubt, the issue should be escalated to your supervisor and/or a call should be made to the authorities.

Eavesdropping

Eavesdropping is a strategy in which a person uses direct observation to "listen" in to a conversation. This could be a person hiding around the corner or a person tapping into a phone conversation. Soundproof rooms are often employed to stop eavesdropping, and encrypted phone sessions can also be implemented.

Baiting

Baiting is a strategy in which a malicious individual leaves malware-infected removable media such as a USB drive or optical disc lying around in plain view. It might have an interesting logo or distinctive look about it. When a person takes it and connects it to his or her computer, the malware infects the computer and attempts to take control of it and/or the network the computer is a member of.



Watering Hole Attack

The ***watering hole attack*** is a strategy that targets users based on the common websites that they frequent. The attacker loads malware beforehand on one or more websites in the hopes that the user will access those sites and activate the malware, ultimately infecting the user's system and possibly spreading through the network. To figure out the browsing habits of users, the attacker might guess or use direct observation. So, this attack may also build on other social engineering methods such as eavesdropping, pretexting, and phishing.

Popular websites such as Google, Microsoft, and so on will be difficult to infect with malware. It's the smaller websites that the attacker will go after. For example, let's consider a company that manufactures widgets. Chances are that the company will need to purchase plastic and other resources to build the widgets. It follows that users will connect to suppliers' websites often via the Internet or possibly an intranet. Typically, suppliers' websites are known for a lack of security and make excellent targets. If many users in the company go to these same websites, and often, it's just a matter of time before one clicks on the wrong website element or gets tricked in another manner. Then malware gets installed to the client computer and possibly spreads throughout the company. An attacker might also redirect users to other websites where more hardcore malware (such as ransomware) or other scams are located.

The problem is that you, as a security administrator, can't actively prevent the malware on the targeted websites. You can suggest prevention methods to those companies—such as software patches and secure coding—but can't force them into action. So, you should focus on localized prevention methods including training users, reducing web browser functionality, blocking known malicious websites based on threat intelligence, and monitoring in the form of antimalware software and IDS/IPS.



Typo Squatting

Typo squatting (or *typosquatting*) is a technique used by adversaries that leverages human error when typing a URL in their web browser. For instance, users often enter typos (or a wrong URL) for a given website. Let's suppose you want to go to

google.com, but instead typed gogle.com. An attacker can register the "gogle.com" domain and impersonate the real website and host malware or perform any other malicious technique to compromise your system.

Key Topic Influence Campaigns, Principles of Social Engineering, and Reasons for Effectiveness

The following are several motivation techniques used by social engineers:

- **Authority:** A social engineer shows confidence and perhaps authority—whether legal, organizational, or social authority.
- **Intimidation:** Attackers can use intimidation to manipulate their victims to perform some action or to reveal sensitive information.
- **Consensus:** Consensus, or "social proof," is a psychological phenomenon in which an individual is not able to determine the appropriate mode of behavior. For example, you might see others acting or doing something in a certain way and might assume that it is appropriate. Social engineers might use this tactic when an individual enters an unfamiliar situation that he or she doesn't know how to deal with. Social engineers might manipulate multiple people at once by using this technique.
- **Scarcity:** It is possible to use scarcity to create a feeling of urgency in a decision-making context. Specific language can be used to heighten urgency and manipulate victims. Salespeople often use scarcity to manipulate clients (for example, telling a customer that an offer is for today only or that there are limited supplies). Social engineers use similar techniques.
- **Familiarity:** Individuals can be influenced by things or people they like or are familiar with. Social engineers take advantage of these human vulnerabilities to manipulate their victims.
- **Trust:** Attackers take advantage of the trust a person has in another person or organization in order to influence them to perform some action or reveal sensitive information.
- **Urgency:** It is possible to manipulate a person with a sense of immediate urgency to prompt him or her to act promptly. Using urgency, social engineers force their victims to act quickly to avoid or rectify a perceived dangerous or painful situation.

TIP Nowadays, threat actors use automated bots in ***social media*** sites like Twitter and Facebook to influence the sentiment of a given user. These bots are used to try to manipulate public sentiment on contentious issues including political events, gun control, abortion, and many more. ***Hybrid warfare*** is a subject that was originally employed by the armed forces. However, attackers nowadays use hybrid warfare techniques in cyber and influential campaigns to manipulate people to believe something that may not be true by using different types of propaganda that are often shared in social media sites.

Key Topic

User Security Awareness Education

The user could be the single most exploitable resource of a company. I think that the previous section about social engineering provides a good argument to support that opinion. People aren't computers. Some might be compared to robots (perhaps unfairly), but no matter how logical and disciplined a person might be, there is always the potential for human error. So, until firewalls are developed for people's brains, education becomes the best method to prevent social engineering attacks and malware infection, not to mention user error.

As an IT manager, I have always made it a point to incorporate technology and security training for as many employees as possible. This would include classroom/conference room training, written materials, digital courses, even funny posters in the cafeteria. I'd use whatever platform I could to get the point across. The key is to make it accessible to the user and to make it interesting and fun.

There are several roadblocks when it comes to user training. The first is organizational acceptance. Are the executives of a company on board with the idea? As time moves on, we see that more and more executives include user education as a matter of course. However, if you come across an individual who is against the idea because of a "lack" of budgeting or time, you can simply show that person a news article about one of the many successful attacks that have occurred recently. Then show a case study of the amount of time and money that the affected company lost due to the attack, and—in most cases—how easily it could have been prevented. Then, there are the employees to be trained themselves. Some will put up a fight when it comes to education. Again, the secret ingredient here is to pique the interest of the users. Get them involved, make it fun, create a reward system, and use your imagination. Some organizations employ IT trainers on a full-time basis or as consultants. Good IT trainers know how to get through to the typical employee. Other organizations will offer incentives for attending training, or penalties for not attending—though statistics show that incentives usually work better than penalties.

You can also attempt role-playing. I'm not talking about a role-playing game such as *Dungeons & Dragons*! Rather, having the employees act out different organizational roles, such as system administrators, privileged users, executive users, data owners, system owners, and of course, typical end users. Throw in the hacker and/or con artist as the "bad guy" role and you can really teach in a fun way. Create scenarios where people can learn in a tangible way how attacks are carried out and how they can be prevented. Quiz the employees, but keep it light. The idea is for employees to expand their knowledge into other areas of the company—a sort of table-top job rotation, so to speak. Develop the situations and solutions properly, and the whole organization becomes a stronger and more secure unit, thanks to you. Come on, you know you always aspired to be a Dungeon Master!

Finally, there is the time factor. People have projects and tasks to complete and usually aren't even given enough time for that! Where does the time for training come from? This is when the mindset of loyalty to the company comes in—and Human Resources can usually be helpful in cultivating that mindset. The whole outlook should be based on the idea of overall efficiency and benefits to the company and individual. By sitting in on security training, the user will save time over the long term and will ultimately become a more knowledgeable person.

Anyway, for the Security+ certification, *how* the training gets accomplished isn't as important as *what* is covered in training. The following is a basic list of rules you can convey when training employees:

Key Topic

- Never, under any circumstances, give out any authentication details such as passwords, PINs, company ID, tokens, smart cards, and so on.
- Always shield keypads and screens when entering authentication information.
- Adhere to the organization's *clean desk policy*, which states that all documents, electronics, personally owned devices, and other items be put away (or locked away) when you are not at your desk, or other work area.
- Always screen your email and phone calls carefully and keep a log of events. This is also known as communications *vetting*.
- Use encryption when possible to protect emails, phone calls, and data.
- If there is any doubt as to the legitimacy of a person, email, or phone call, document the situation and escalate it to your supervisor, security, or the authorities.
- Never pick up, and make use of, any unknown removable media.
- Always shred any sensitive information destined for the garbage or recycling.
- Always comply with company policy when it comes to data handling and disposal. For example, if a hard drive, USB flash drive, memory stick, or optical

disc is no longer being used, make sure it is disposed of properly. If you are not sure, contact the IT department or facilities department of the organization to find out if it should be recycled, or destroyed.

- Always track and expedite shipments.
- Be extremely careful when using a web browser. Double-check everything that is typed before pressing Enter or clicking Go. Don't click on anything unless you know exactly what it is. Web-based attacks account for a huge percentage of damage to organizations.

When training employees, try to keep them interested; infuse some fun and be silly if you want to. For instance, the first bullet said to *never* give out authentication details. Pundits tell us to never say never. Well, if it's okay for them, then it's okay for us. And in this case, it's vital to the health of your organization—not to mention you. You see what I mean? Or, if you don't want to be silly, then consider imparting some real examples. Use examples of social engineering so that your trainees can make the connection between actual social engineering methods and their defenses. Make them understand that social engineers don't care how powerful an organization's firewall is or how many armed guards the company has. They get past technology and other types of security by exploiting the weaknesses inherent in human nature.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 1-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 1-2 Key Topics for Chapter 1

Key Topic Element	Description	Page Number
Section	Phishing and Spear Phishing	9
Section	Smishing	12
Section	Vishing	12
Section	Spam and Spam over Internet Messaging (SPIM)	13

Key Topic Element	Description	Page Number
Section	Dumpster Diving	13
Section	Shoulder Surfing	14
Section	Pharming	14
Section	Piggybacking and Tailgaiting	15
Section	Eliciting Information	15
Section	Whaling	16
Section	Prepending	17
Section	Identity Fraud	17
Section	Invoice Scams	17
Section	Credential Harvesting	18
Section	Reconnaissance	18
Section	Hoaxes	19
Section	Impersonation or Pretexting	19
Section	Watering Hole Attack	20
Section	Typo Squatting	20
Section	Influence Campaigns, Principles of Social Engineering, and Reasons for Effectiveness	21
Section	User Security Awareness Education	22
List	Rules you can convey when training employees	23

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

phishing, pretexting, spear phishing, whaling, smishing, vishing, spam, dumpster diving, shoulder surfing, pharming, piggybacking, tailgating, eliciting information, prepending, credential harvesting, hoax, watering hole attack, typo squatting, social media, hybrid warfare

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. You go out the back door of your building and notice someone looking through your company's trash. If this person were trying to acquire sensitive information, what would this attack be known as?
2. User education can help to defend against which type of attacks?
3. You can often configure your email servers or email cloud services to add a message in the email subject line to identify emails that are coming from outside of the organization. This technique is also known as _____.
4. What is the most common reason that social engineering succeeds?
5. In which two environments would social engineering attacks be most effective?
6. What is a social engineering technique used by adversaries that leverages user errors ("typos") when entering a given URL in their web browser for a given website?
7. A man pretending to be a data communications repair technician enters your building and states that there is networking trouble and he needs access to the server room. What is this an example of?
8. Turnstiles, double entry doors, and security guards are all preventive measures for what kind of social engineering?
9. What is the social engineering technique where the attacker redirects the victim from a valid website or resource to a malicious one that could be made to appear as the valid site to the user?
10. Why would you implement password masking and privacy screens/filters?

This page intentionally left blank



This chapter covers the following topics related to Objective 1.2 (Given a scenario, analyze potential indicators to determine the type of attack) of the CompTIA Security+ SY0-601 certification exam:

- Malicious Software (Malware)
  - Ransomware
  - Trojans
  - Worms
  - Potentially unwanted programs (PUPs)
  - Fileless virus
  - Command and control
  - Bots
  - Cryptomalware
  - Logic bombs
  - Spyware
  - Keyloggers
  - Remote access Trojan (RAT)
  - Rootkit
  - Backdoor
- Password attacks
  - Spraying
  - Dictionary
  - Brute force (Offline and Online)
  - Rainbow tables
  - Plaintext/unencrypted
- Physical attacks
  - Malicious universal serial bus (USB) cable
  - Malicious flash drive
  - Card cloning
  - Skimming
- Adversarial artificial intelligence (AI)
  - Tainted training data for machine learning (ML)
  - Security of machine learning algorithms
- Supply-chain attacks
- Cloud-based vs. on-premises attacks
- Cryptographic attacks
  - Birthday
  - Collision
  - Downgrade

CHAPTER 2

Analyzing Potential Indicators to Determine the Type of Attack

To combat the various security threats that can occur on a computer system, you first need to classify them. Then you need to define how these threats can be delivered to the target computer. Afterward you can discuss how to prevent security threats from happening and troubleshoot them if they do occur. In this chapter, let's start with the most common computer threat and probably the deadliest—malicious software.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 2-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 2-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Malicious Software (Malware)	1–5
Password Attacks	6–7
Physical Attacks	8
Adversarial Artificial Intelligence (AI)	9
Supply-Chain Attacks	10
Cloud-based vs. On-premises Attacks	11
Cryptographic Attacks	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. You logged in to your laptop and noticed a message saying that all your files have been encrypted and to recover them you need to pay $1,000 in Bitcoins. What has your system been infected with?
   a. Ransomware
   b. Worm
   c. Keylogger
   d. None of these answers are correct.

2. What type of malware can look like legitimate software but then performs negative actions to manipulate your system?
   a. Trojan
   b. Ransomware
   c. Worm
   d. None of these answers are correct.

3. Which malware type can allow an attacker to gain administrator privileges?
   a. Keylogger
   b. Rootkit
   c. Ransomware
   d. All of these answers are correct.

4. What type of malware does not reside on the hard drive of a computer?
   a. Ransomware
   b. Botnets
   c. Fileless malware
   d. None of these answers are correct.

5. Bots in a botnet typically receive instructions from which of the following?
   a. A command-and-control (C2) server
   b. A zombie system
   c. A malvertising site
   d. All of these answers are correct.

6. An attacker using John the Ripper, which uses a wordlist, is an example of which of the following?
   a. Social engineering attack
   b. Dictionary password attack
   c. Buffer overflow attack
   d. Cross-site request forgery attack

7. An attacker using a large number of usernames with a few commonly used passwords is considered what kind of attack?
   a. Password spraying
   b. Credential harvesting
   c. Password cracking
   d. None of these answers are correct.

8. What type of attack occurs when an attacker captures credit card information or information from other similar cards (gift cards, loyalty cards, identification cards, and so on)?
   a. Skimming
   b. Shimming
   c. SIM cloning
   d. None of these answers are correct.

9. Which of the following techniques are used to attack machine learning (ML) implementations?
   a. Tainting of data to cause errors in the outcome of the ML solution
   b. Overfitting attacks
   c. ML transfer attacks
   d. All of these answers are correct.

10. You purchased a brand-new Internet of Things (IoT) device and noticed that it started collecting personal information (PI) and attempted to send your data by communicating with random IP addresses. You noticed that an implant could have been installed during the manufacturing of the product. What type of attack might this be?
    a. Supply-chain attack
    b. Cross-site scripting
    c. Return to libc attack
    d. Masquerading attack

11. Which of the following attacks are made against cloud implementations?
    a. API attacks
    b. DNS attacks
    c. VM escape attacks
    d. All of these answers are correct.

12. An attacker attempts to force an application to roll back the version of TLS (from TLS version 1.3 to 1.0). What is the name of this type of attack?
    a. Privilege escalation
    b. Downgrade attack
    c. Cracking
    d. Fuzzing

Foundation Topics

Malicious Software (Malware)

Key Topic

Malicious software, or ***malware***, is software designed to infiltrate a computer system and possibly damage it without the user's knowledge or consent. *Malware* is a broad term used by computer professionals to include viruses, worms, Trojan horses, spyware, rootkits, adware, and other types of undesirable software.

Of course, as a security professional, you don't want malware to infect your systems, but to defend against it, you first need to define it and categorize it. Then you can put preventive measures into place. It's also important to locate and remove or quarantine malware from a computer system in the case that it does manifest itself.

For the exam, you need to know about several types of malware. Over the past several years, an emphasis shift from viruses to other types of malware, such as spyware and ransomware, has occurred. Most people know about viruses and have some kind of antivirus software running. However, many people are still confused about the various other types of malware, how they occur, and how to protect against them. As a result, computer professionals spend a lot of time fixing malware issues (that are not virus related) and training users on how to protect against them in the future. However, viruses are still a valid foe, so let's start by discussing them.

A computer virus is code that runs on a computer without the user's knowledge; it infects the computer when the code is accessed and executed. For viruses to do their dirty work, they first need to be executed by the user in some way. A virus also has reproductive capability and can spread copies of itself throughout the computer—if it is first executed, by the user or otherwise. By infecting files accessed by other computers, the virus can spread to those other systems as well. The problem is that computers can't call in sick on Monday; they need to be up and running as much as possible, more than the average human.

Ransomware and Cryptomalware

Key Topic

Some less than reputable persons use a particularly devious malware known as ***ransomware***—a type of malware that restricts access to a computer system and demands that a ransom be paid. Ransomware informs the user that in order to decrypt the files or unlock the computer to regain access to the files, a payment would have to be made to one of several banking services (typically crypto currencies like Bitcoin). It often propagates as a Trojan or worm, and usually makes use of encryption to cause the user's files to be inaccessible. This use of encryption is also known as cryptoviral extortion. Examples of ransomware include WannaCry, NotPetya, Nyetya, SamSam, BadRabbit, and others. One of the most common infection

methods of ransomware is via phishing and spear phishing attacks. You learned about phishing and spear phishing in Chapter 1, "Comparing and Contrasting Different Types of Social Engineering Techniques."

Key Topic

Ransomware is also often referred to as ***cryptomalware***. Today's cryptomalware leverages advanced encryption techniques to prevent files from being decrypted without a unique key.

Figure 2-1 shows a message displayed to the victim of a ransomware attack. The ransomware illustrated in this figure is the well-known ransomware called WannaCry.

Key Topic



FIGURE 2-1 A Ransomware Attack

NOTE Sometimes a user will inadvertently access a fraudulent website (or pop-up site) saying that all the user's files have been encrypted and payment is required to decrypt them; some imposing government-like logo will accompany the statement. But many of these sites don't actually encrypt the user's files. In this case, we are talking about plain old extortion with no real damage done to the computer or files. These types of sites can be blocked by pop-up blockers, phishing filters, and the user's common sense when clicking searched-for links.

Trojans

Key Topic

Trojan horses, or simply ***Trojans***, appear to perform desirable functions but are actually performing malicious functions behind the scenes. They are not technically viruses and can easily be downloaded without being noticed. They can also be transferred to a computer by way of removable media, especially USB flash drives. One example of a Trojan is a file that is contained within a downloaded program such as a key generator—known as a keygen used with pirated software—or another executable. If users complain about slow system performance and numerous antivirus alerts, and they recently installed a questionable program from the Internet or from a USB flash drive, their computers could be infected by a Trojan.

Remote Access Trojans (RATs) and Rootkits

Key Topic

Remote access Trojans (RATs) are the most common type of Trojan. Examples include Back Orifice, NetBus, and SubSeven. Their capability to allow an attacker higher administration privileges than those of the owner of the system makes them quite dangerous. The software effectively acts as a remote administration tool, which happens to be another name for the *RAT* acronym. These programs have the capability to scan for unprotected hosts and make all kinds of changes to a host when connected. They are not *necessarily* designed to be used maliciously but are easy for an average person to download and manipulate computers with. Worse, when a target computer is controlled by an attacker, it could easily become a robot, or simply a *bot*, carrying out the plans of the attacker on command. We discuss bots later in this chapter.

RATs can also be coded in many programming languages (such as PHP, Ruby, and Python) to allow remote access to websites. An example of this is the web shell, which has many permutations. It allows an attacker to remotely configure a web server without the user's consent. Quite often, the attacker will have cracked the FTP password in order to upload the RAT.

RATs are often used to persistently target a specific entity such as a government or a specific corporation. One example is the PlugX RAT. Malicious software such as this is known as an advanced persistent threat (APT). Groups that have vast resources at their disposal might make use of these APTs to carry out objectives against large-scale adversaries. APT groups could take the form of large hacker factions, and even some corporations and governments around the globe.

Key Topic

A ***rootkit*** is a type of software designed to gain administrator-level control over a computer system without being detected. The term is a combination of the words *root* (meaning the root user in a UNIX/Linux system or administrator in a Windows system) and *kit* (meaning software kit). Usually, the purpose of a rootkit is to perform malicious operations on a target computer at a later date without the knowledge of the administrators or users of that computer. A rootkit is a variation

on the virus that attempts to dig in to the lower levels of the operating system—components of the OS that start up before any antimalware services come into play. Rootkits can target the UEFI/BIOS, boot loader, kernel, and more. An example of a boot loader rootkit is the Evil Maid Attack, which can extract the encryption keys of a full disk encryption system, as discussed later. Another (more current) example is the Alureon rootkit, which affects the master boot record (MBR) and low-level system drivers (such as atapi.sys). This particular rootkit was distributed by a botnet and affected over 200,000 (known) Microsoft operating systems.

Rootkits are difficult to detect because they are activated before the operating system has fully booted. A rootkit might install hidden files, hidden processes, and hidden user accounts. Because rootkits can be installed in hardware or software, they can intercept data from network connections, keyboards, and so on.

A successfully installed rootkit enables unauthorized users to gain access to a system and act as the root or administrator user. Rootkits are copied to a computer as a binary file. This binary file can be detected by signature-based and heuristic-based antivirus programs; however, after the rootkit is executed, it can be difficult to detect. The reason is that most rootkits are collections of programs working together that can make many modifications to the system. When subversion of the operating system takes place, the OS can't be trusted, and it is difficult to tell if your antivirus programs run properly or if any of your other efforts have any effect. Although security software manufacturers are attempting to detect running rootkits, it is doubtful that they will be successful. The best way to identify a rootkit is to use removable media (a USB flash drive or a special rescue disc) to boot the computer. This way, the operating system is not running, and therefore, the rootkit is not running, making it much easier to detect by the external media.

Sometimes, rootkits will hide in the MBR. Often, operating system manufacturers recommend scrubbing the MBR (rewriting it, for example, within System Recovery Options or another recovery environment) and then scanning with antivirus software. This solution depends on the type of rootkit. The use of GPT in lieu of MBR helps to discourage rootkits. I suggest using GPT whenever possible.

Unfortunately, because of the difficulty involved in removing a rootkit, the best way to combat rootkits is to reinstall all software (or reimage the system). Generally, an IT technician, upon detecting a rootkit, will do just that because reinstalling usually takes less time than attempting to fix all the rootkit issues; plus, it can verify that the rootkit has been removed completely.

Worms



A ***worm*** is much like a virus except that it self-replicates, whereas a virus does not. It does this in an attempt to spread to other computers. Worms take advantage of

security holes in operating systems and applications, including backdoors, which we discuss later. They look for other systems on the network or through the Internet that are running the same applications and replicate to those other systems. With worms, the user doesn't need to access and execute the malware. A virus needs some sort of carrier to get it where it wants to go and needs explicit instructions to be executed, or it must be executed by the user. The worm does not need this carrier or explicit instructions to be executed.

Going back in history, a well-known example of a worm is Nimda (*admin* backward), which propagated automatically through the Internet in 22 minutes in 2001, causing widespread damage. It spread through network shares, mass emailing, and operating system vulnerabilities.

Sometimes, the worm does not carry a payload, meaning that in and of itself, it does not contain code that can harm a computer. It may or may not include other malware, but even if it doesn't, it can cause general disruption of network traffic and computer operations because of the very nature of its self-replicating abilities.

In the late 2010s, different types of ransomware (including WannaCry) propagated like a worm by scanning and infecting vulnerable systems leveraging a known Windows SMB vulnerability.

Fileless Virus



Malware doesn't have to reside on the hard drive of a computer. It can also reside within RAM (and possibly other locations). ***Fileless malware***—also known as non-malware—functions without putting malicious executables within the file system and instead works in a memory-based environment. Sound potent? It can be, but the attacker will usually need to have remote access to the system—via SSH, a RAT, or otherwise—in order to deploy the malware. The term *fileless* is somewhat of a misnomer because the attack may actually contain files. The systems affected, such as ATMs, often suffer from weak endpoint security, and the organization that owns those systems might also have ineffective access control policies.

Command and Control, Bots, and Botnets



I know what you're thinking: the names of these attacks and delivery methods are getting a bit ridiculous. But bear with me; they make sense and are deadly serious. Allow me to explain—malware can be distributed throughout the Internet by a group of compromised computers (***bots***), known as a ***botnet***, and controlled by a master computer (***command-and-control*** [C2] server). The individual compromised computers in the botnet are called bots and also often called zombies because they are unaware of the malware that has been installed on them. These compromises can occur in several ways, including automated distribution of the malware from one

zombie computer to another. Now imagine if all the zombie computers had a specific virus or other type of attack loaded, and a logic bomb (defined a bit later) was also installed, ready to set off the malware at a specific time. If this were done to hundreds or thousands of computers, a synchronized attack of great proportions could be enacted on just about any target. Often, this is known as a distributed denial-of-service (DDoS) attack and is usually perpetuated on a particularly popular server, one that serves many requests. If a computer on your network is continually scanning other systems on the network, is communicating with an unknown server, and/or has hundreds of outbound connections to various websites, chances are the computer is part of a botnet. But botnets can be used for more than just taking down a single target. They can also be used to fraudulently obtain wealth and even crypto mining.

Figure 2-2 shows a botnet controlled by a command-and-control server. The attacker (via the C2) sends instructions to the bots in the botnet. A botnet could be composed of any compromised system. These could be user endpoints, as well as Internet of Things (IoT) devices (such as cameras, sensors, and industrial control systems).

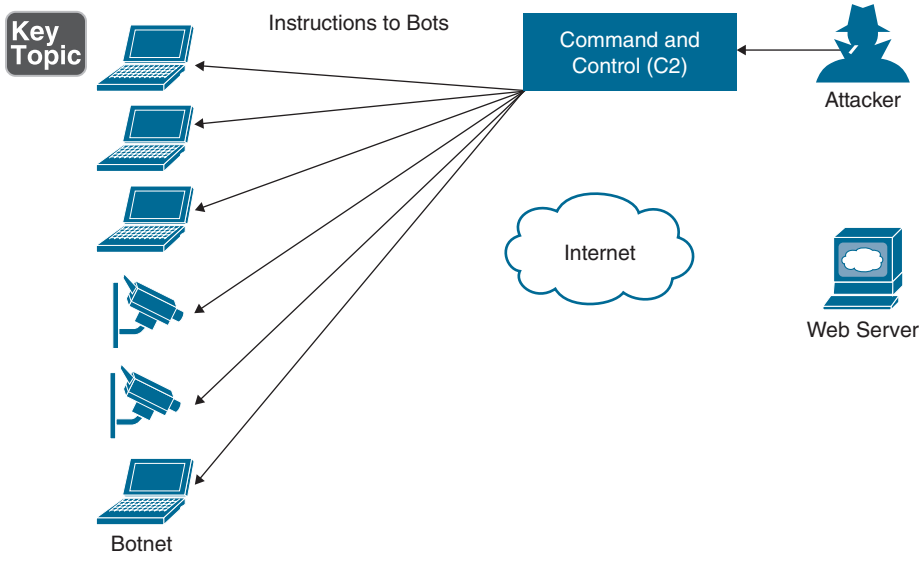


FIGURE 2-2 A Botnet and C2 Example

After the bots receive the instructions from the C2, they perform a given action. In the example illustrated in Figure 2-3, the bots start sending numerous IP packets to a web server (the victim) to perform a DDoS attack. The purpose of this attack is to consume system resources and prevent authorized users from accessing the web server. You will learn additional details about DDoS attacks in Chapter 4, "Analyzing Potential Indicators Associated with Network Attacks."

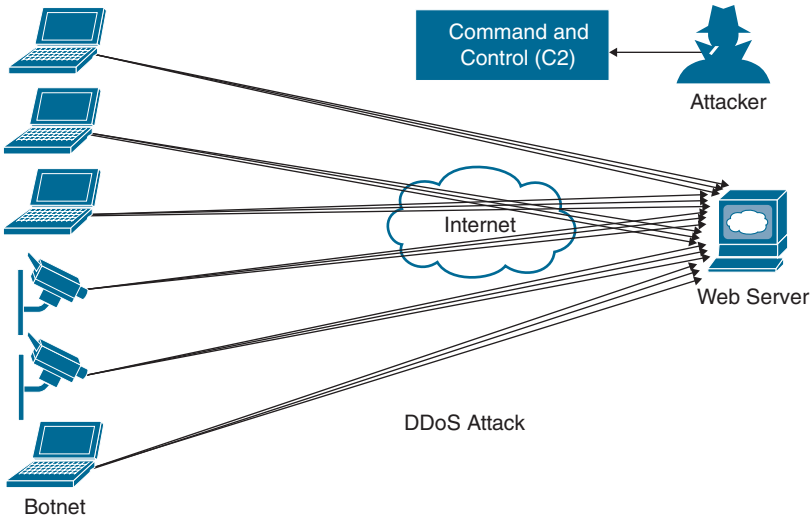


FIGURE 2-3 The Botnet Performing a DDoS Attack

Logic Bombs

Key Topic

A ***logic bomb*** is code that has, in some way, been inserted into software; it is meant to initiate one of many types of malicious functions when specific criteria are met. Logic bombs blur the line between malware and a malware delivery system. They are indeed unwanted software but are intended to activate viruses, worms, or Trojans at a specific time. Trojans set off on a certain date are also referred to as *time bombs*. The logic bomb ticks away until the correct time, date, and other parameters have been met. So, some of the worst bombs do not incorporate any explosion whatsoever. The logic bomb could be contained within a virus or loaded separately. Logic bombs are more common in the movies than they are in real life, but they do happen—and with grave consequences. But more often than not, they are detected before they are set off. If you, as a security administrator, suspect that you have found a logic bomb or a portion of the code of a logic bomb, you should notify your superior immediately and check your organization's policies to see if you should take any other actions. Action could include placing network disaster recovery processes on standby, notifying the software vendor, and closely managing usage of the software, including, perhaps, withdrawing it from service until the threat is mitigated. Logic bombs are the evil cousin of the Easter egg.

Easter eggs historically have been platonic extras added to an OS or application as a sort of joke; often, they were missed by quality control and subsequently released by the manufacturer of the software. Easter eggs are not normally documented (being tossed in last minute by humorous programmers) and are meant to be harmless, but

nowadays they are not allowed by responsible software companies and are thoroughly scanned for. Because an Easter egg (and who knows what else) can possibly slip past quality control, and because of the growing concerns about malware in general, many companies have adopted the idea of *trustworthy computing*, which is a newer concept that sets standards for how software is designed, coded, and checked for quality control. Sadly, as far as software goes, the Easter egg's day has passed.

Potentially Unwanted Programs (PUPs) and Spyware

Key Topic

Spyware is a type of malicious software either downloaded unwittingly from a website or installed along with some other third-party software. Usually, this malware collects information about the user without the user's consent. Spyware could be as simple as a piece of code that logs what websites you access, or it could go as far as a program that records your keystrokes (known as a keylogger). Spyware is also associated with advertising (those pop-ups that just won't go away!) and is sometimes related to malicious advertising, or *malvertising*—the use of Internet-based advertising (legitimate and illegitimate) to distribute malicious software.

Spyware can possibly change the computer configuration without any user interaction—for example, redirecting a browser to access websites other than those wanted. One example (of many) of spyware is the Internet Optimizer, which redirects Internet Explorer error pages out to other websites' advertising pages. Spyware can even be taken to the next level and be coded in such a way to hijack a user's computer. Going beyond this, spyware can be used for cyber espionage, as was the case with Red October, which was installed to users' computers when they unwittingly were redirected to websites with special PHP pages that exploited a Java flaw, causing the download of the malware.

Key Topic

Adware usually falls into the realm of spyware because it pops up advertisements based on what it has learned from spying on the user. *Grayware* is another general term that describes applications that are behaving improperly but without serious consequences. It is associated with spyware, adware, and joke programs. Very funny… not. Adware and grayware are types of ***potentially unwanted programs (PUPs)***.

Preventing spyware and PUPs works in much the same manner as preventing other types of malware when it comes to updating the operating system and using a firewall. Also, because spyware is as common as viruses, antivirus companies and OS manufacturers add antispyware components to their software. Here are a few more things you can do to protect your computer in the hopes of preventing spyware:

- Use (or download) and update built-in antispyware programs such as Windows Defender. Be sure to keep the antispyware software updated.
- Adjust web browser security settings. For example, disable (or limit) cookies, create and configure trusted zones, turn on phishing filters, restrict unwanted websites, turn on automatic website checking, disable scripting (such as

JavaScript and ActiveX), and have the browser clear all cache on exit. All of these things can help filter out fraudulent online requests for usernames, passwords, and credit card information, which is also known as web-page spoofing. Higher security settings can also help fend off session hijacking, which is the act of taking control of a user session after obtaining or generating an authentication ID.

- Uninstall unnecessary applications and turn off superfluous services (for example, Remote Desktop services or FTP if they are not used).
- Educate users on how to surf the web safely. User education is actually the number one method of preventing malware! Access only sites believed to be safe, and download only programs from reputable websites. Don't click OK or Agree to close a window; instead press Alt+F4 on the keyboard to close that window, or use the Task Manager to close out of applications. Be wary of file-sharing websites and the content stored on those sites. Be careful of emails with links to downloadable software that could be malicious.
- Consider technologies that discourage spyware. For example, use a browser that is less susceptible to spyware, and consider using virtual machines.
- Verify the security of sites you visit by checking the certificate or by simply looking for HTTPS in the URL.

Here are some common symptoms of spyware:

- The web browser's default home page has been modified.
- A particular website comes up every time you perform a search.
- Excessive pop-up windows appear.
- The network adapter's activity LED blinks frequently when the computer shouldn't be transmitting data.
- The firewall and antivirus programs turn off automatically.
- New programs, icons, and favorites appear.
- Odd problems occur within windows (slow system, applications behaving strangely, and such).
- The Java console appears randomly.

To troubleshoot and repair systems infected with spyware, first disconnect the system from the Internet (or simply from the local-area network). Then try uninstalling the program from the Control Panel or Settings area of the operating system. Some of the less malicious spyware programs can be fully uninstalled without any

residual damage. Be sure to reboot the computer afterward and verify that the spyware was actually uninstalled! Next, scan your system with the AV software to remove any viruses that might have infested the system, which might get in the way of a successful spyware removal. In Windows, do this in the recovery environment (for example, Safe Mode) if the AV software offers that option.

Keyloggers

Key Topic

An attacker may use a ***keylogger*** to capture every keystroke of a user in a system and steal sensitive data (including credentials). There are two main types of keyloggers: keylogging hardware devices and keylogging software. A hardware (physical) keylogger is usually a small device that can be placed between a user's keyboard and the main system. Software keyloggers are dedicated programs designed to track and log user keystrokes.

TIP Keyloggers are legal in some countries and designed to allow employers to oversee the use of their computers. However, recent regulations like the General Data Protection Regulation (GDPR) in the European Union (EU) have made keyloggers a very sensitive and controversial topic. Threat actors use keyloggers for the purpose of stealing passwords and other confidential information.

There are several categories of software-based keyloggers:

- **Kernel-based keylogger:** With this type of keylogger, a program on the machine obtains root access to hide itself in the operating system and intercepts keystrokes that pass through the kernel. This method is difficult both to write and to combat. Such keyloggers reside at the kernel level, which makes them difficult to detect, especially for user-mode applications that don't have root access. They are frequently implemented as rootkits that subvert the operating system kernel to gain unauthorized access to the hardware. This makes them very powerful. A keylogger using this method can act as a keyboard device driver, for example, and thus gain access to any information typed on the keyboard as it goes to the operating system.
- **API-based keylogger:** With this type of keylogger, compromising APIs reside inside a running application. Different types of malware have taken advantage of Windows APIs, such as GetAsyncKeyState() and GetForeground Window(), to perform keylogging activities.
- **Hypervisor-based keylogger:** This type of keylogger is effective in virtual environments, where the hypervisor could be compromised to capture sensitive information.

- **Web form–grabbing keylogger:** Keyloggers can steal data from web form submissions by recording the web browsing on submit events.
- **JavaScript-based keylogger:** Malicious JavaScript tags can be injected into a web application and then capture key events (for example, the onKeyUp() JavaScript function).
- **Memory-injection-based keylogger:** This type of keylogger tampers with the memory tables associated with the browser and other system functions.

Backdoors

Key Topic

Backdoors are pieces of software, malware, or configuration changes that allow attackers to control a victim's system remotely. For example, a backdoor can open a network port on the affected system so that the attacker can connect and control the system.

A backdoor can be installed by an attacker to either allow future access or to collect information to use in further attacks. When the threat actor gains access to a system, the attacker usually wants future access so he or she can control the system and can carry out other attacks. The attacker wants that access to be really easy. Many backdoors are installed by users clicking something without realizing that the link that they clicked or the file that they opened is actually a threat. A backdoor can be implemented as a result of malware, a virus, or a worm.

Malware Delivery Mechanisms

Malware is not sentient (…not yet) and can't just appear out of thin air; it needs to be transported and delivered to a computer or installed on a computer system in some manner. This can be done in several ways. The simplest way would be for attackers to gain physical access to an unprotected computer and perform their malicious work locally. But because obtaining physical access can be difficult, this can be done in other ways, as shown in the upcoming sections. Attackers can use some of these methods to simply gain access to a computer, make modifications, and so on, in addition to delivering the malware.

The method that a threat uses to access a target is known as a threat vector. Collectively, the means by which an attacker gains access to a computer in order to deliver malicious software is known as an attack vector. Probably the most common attack vector is via software.

Malware can be delivered via software in many ways. A person who emails a zipped file might not even know that malware also exists in that file. The recipients of the email will have no idea that the extra malware exists unless they have software

to scan their email attachments for it. Malware could also be delivered via FTP. Because FTP servers are inherently insecure, it's easier than you might think to upload insidious files and other software. Malware is often found among peer-to-peer (P2P) networks and bit torrents. Great care should be taken by users who use these technologies. Malware can also be embedded within, and distributed by, websites through the use of corrupting code or bad downloads. Malware can even be distributed by advertisements. And of course, removable media can victimize a computer as well. Optical discs, USB flash drives, memory cards, and connected devices such as smartphones can easily be manipulated to automatically run malware when they are inserted into the computer. (This is when AutoPlay/AutoRun is not your friend!) The removable media could also have hidden viruses or worms and possibly logic bombs (discussed earlier) configured to set that malware off at specific times.

Potential attackers also rely on user error. For example, if a user is attempting to access a website but types the incorrect domain name by mistake, that user could be redirected to an altogether unwanted website, possibly malicious in nature. This type of attack is known as typo squatting or URL hijacking.

NOTE You learned details about typo squatting in Chapter 1.

URL stands for Uniform Resource Locator, which is the web address that begins with http or https. The potential attacker counts on the fact that millions of typos are performed in web browsers every day. These attackers "squat" on similar (but not exact) domain names. Once the user is at the new and incorrect site, the system becomes an easy target for spyware and other forms of malware. Some browsers come with built-in security such as antiphishing tools and the capability to auto-check websites that are entered, but the best way to protect against this issue is to train users to be careful when typing domain names.

Another way to propagate malware is to employ automation and web attack-based software "kits," also known as exploit kits. Exploit kits are designed to run on web servers and are usually found in the form of PHP scripts. They target client computers' software vulnerabilities that often exist within web browsers. One example of an exploit kit is the Blackhole exploit kit. It is used (and purchased) by potential attackers to distribute malware to computers that meet particular criteria, while the entire process is logged and documented.

NOTE The automating of cybercrime, and the software used to do so, is collectively referred to as *crimeware*.

Active interception normally includes a computer placed between the sender and the receiver to capture and possibly modify information. If a person can eavesdrop on your computer's data session, then that data can be stolen, modified, or exploited in other ways. Examples include session theft and man-in-the-middle (MITM) attacks. For more information on these attacks, see Chapter 4.

You Can't Save Every Computer from Malware!

On a sad note, sometimes computers become so infected with malware that they cannot be saved. In this case, the data should be backed up (if necessary by removing the hard drive and slaving it to another system), and the operating system and applications reinstalled. The UEFI/BIOS of the computer should also be flashed. After the reinstall, the system should be thoroughly checked to make sure that there are no residual effects and that the system's hard drive performs properly.

Password Attacks

There are many different ways that attackers can steal passwords, crack passwords, and perform other credential-based (password-based) attacks. The following sections describe the most common password-based attacks.

Dictionary-based and Brute-force Attacks

Key Topic

A ***dictionary-based attack*** pulls words from the dictionary or word lists to attempt to discover a user's password. A dictionary attack uses a predefined dictionary to look for a match between the encrypted password and the encrypted dictionary word. Many times, a dictionary attack will recover a user's password in a short period of time if simple dictionary words are used.

A hybrid attack uses a dictionary or a word list and then prepends and appends characters and numbers to dictionary words in an attempt to crack the user's password. These programs are comparatively smart because they can manipulate a word and use its variations. For example, take the word *password*. A hybrid password audit would attempt variations such as *1password*, *password1*, *p@ssword*, *pa44w0rd*, and so on. Hybrid attacks might add some time to the password-cracking process, but they increase the odds of successfully cracking an ordinary word that has had some variation added to it.

A ***brute-force attack*** uses random numbers and characters to crack a user's password. A brute-force attack on an encrypted password can take hours, days, months, or years, depending on the complexity and length of the password. The speed of success depends on the speed of the CPU's power. Brute-force audits attempt every combination of letters, numbers, and characters.

TIP Tools such as 0phtCrack, LCP, Cain and Abel, and John the Ripper can all perform dictionary, hybrid, and brute-force password cracking.

Password Spraying

Key Topic

Password spraying is a type of attack in which an attacker attempts to compromise a system using a large number of usernames with a few commonly used passwords. Traditional brute-force attacks attempt to gain unauthorized access to a single account by guessing the password. This can quickly result in the targeted account getting locked out because commonly used account-lockout policies allow for a limited number of failed attempts (typically three to five) during a set period of time. During a password-spraying attack (also known as the "low-and-slow" method), the malicious actor attempts a single commonly used password (such as *Password1*, *COVID19*, or *asdasd*) against many accounts before moving on to attempt a second password, and so on. This technique allows the actor to remain undetected by avoiding rapid or frequent account lockouts.

Offline and Online Password Cracking

Key Topic

Weak passwords are the bane of today's operating systems and networks. This could be because no policy for passwords was defined, and people naturally gravitate toward weaker, easier-to-remember passwords. Or it could be that a policy was defined but is not complex enough or is out of date. Whatever the reason, it would be wise to scan computers and other devices for weak passwords with an ***online*** or ***offline password cracker***, which uses comparative analysis to break passwords and systematically guesses until it cracks the password. And, of course, a variety of password-cracking programs can help with this. Examples are John the Ripper and hashcat. These programs have a bit of a learning curve but are quite powerful. They can be used to crack all kinds of different passwords on a local system or on remote devices and computers. They sniff out other hosts on the network the way a protocol analyzer would. They are excellent tools to find out whether weak passwords are on the network or to help if users forget their passwords (when password resets are not possible). Attackers can discover hashed passwords on a compromised Windows, Linux, or macOS system and then crack the password on the same system (online) or on a separate dedicated system (offline). The goal is to obtain the original plaintext version of the password.

Figure 2-4 shows how an attacker steals the password hash of a compromised system and then uses that hash to crack the password "offline" in another system with superior compute resources.

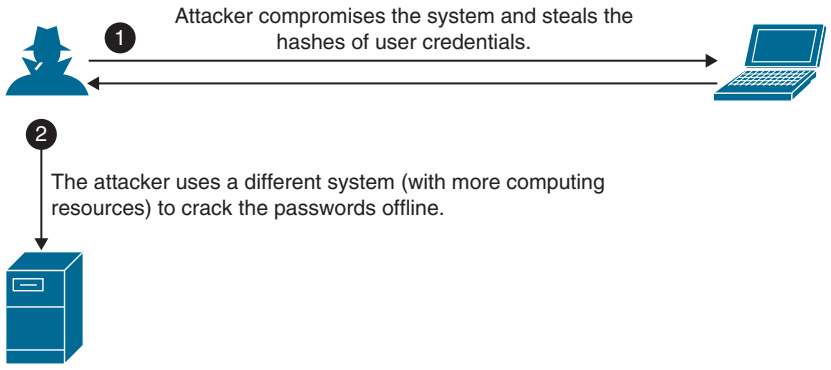


FIGURE 2-4 Offline Password Cracking

Rainbow Tables

Key Topic

A ***rainbow table*** is a precomputed table used to reverse engineer a cryptographic hash function and crack passwords. Attackers can perform cryptanalysis attacks a lot easier using these rainbow tables (as a form of a lookup table). Rainbow tables can be used with the L0phtCrack (https://sectools.org/tool/l0phtcrack/) and RainbowCrack applications.

NOTE This attack can be defeated by implementing salting, which is the randomization of the hashing process.

Plaintext/Unencrypted

Key Topic

If your attempts to guess passwords have not been successful, sniffing or keystroke loggers might offer hope. Think about how much traffic passes over a typical network every day. Most networks handle a ton of traffic, and a large portion of it might not even be encrypted. Password sniffing requires that you have physical or logical access to the device. If that can be achieved, you can sniff the credentials right off the wire as users log in.

One such technique is to pass the hash. Passing the hash enables the hacker to authenticate to a remote server by using the underlying NT LAN Manager (NTLM) and/or LAN Manager (LM) hash of a user's password, instead of using the associated plaintext password. ***Plaintext*** is the term used to describe unencrypted data (in this case an unencrypted password). Mimikatz is a pass the hash application that enables an attacker to authenticate to a remote server using the LM/NTLM hash of a user's password, eliminating the need to crack/brute-force the hashes to

obtain the clear-text password. Because Windows does not salt passwords, they remain static in the Local Security Authority Subsystem Service (LSASS) from session to session until the password is changed. If the password is stored in LSASS and the attacker can obtain a password hash, it can be functionally equivalent to obtaining the clear-text password. Rather than attempting to crack the hash, attackers can simply replay them to gain unauthorized access. You can download this pass the hash toolkit at https://github.com/gentilkiwi/mimikatz.

Physical Attacks

An organization's building is one of its greatest assets and as such it should be properly protected. The following sections detail door access, biometric readers, access logs, and video surveillance to teach you some of the ways to protect the building, its contents, and its inhabitants and to ensure proper authentication when a person enters a building.

Malicious Flash Drives

Key Topic

Malware can be transferred to a computer by way of removable media, especially USB ***malicious flash drives***. For example, an attacker could install a Trojan or ransomware by using a malicious flash drive, if he or she has physical access to the targeted system. Alternatively, the attacker could place the USB flash drive somewhere and use some aspects of social engineering to fool the user into inserting it in his or her system and getting infected. For example, the attacker may attach a key chain or pictures of a dog or a cat. This approach makes the attack a little more personal. The victim may try to find out who those keys and USB flash drive belong to. When he or she inserts the USB flash into a machine, the system is compromised.

Malicious Universal Serial Bus (USB) Cables

Key Topic

In a similar way to malicious USB flash drives, attackers can use ***malicious USB cables*** to compromise systems. Different USB cables are designed to infect connected devices with malware. These malicious USB cables work by injecting keystrokes onto the victim's system when plugged into a USB-capable device.

Card Cloning Attacks

Key Topic

Attackers can perform different ***card cloning attacks***. For example, an attacker can clone a credit card, a smartphone SIM card, or even the badges/cards used to access a building. Specialized software and hardware can be used to perform these cloning attacks.

For instance, a traditional attack on smartphones is SIM cloning (also known as phone cloning), which allows two phones to utilize the same service and allows an attacker to gain access to all phone data. V1 SIM cards had a weak algorithm that made SIM cloning possible (with some expertise). However, V2 cards and higher are much more difficult (if not impossible) to clone due to a stronger algorithm on the chip. Users and administrators should be aware of the version of SIM card being used and update it (or the entire smartphone) if necessary.

There are techniques available to unlock a smartphone from its carrier. Users should be advised against this, and as a security administrator, you should create and implement policies that make unlocking the SIM card difficult if not impossible. Unlocking the phone—making it SIM-free—effectively takes it off the grid and makes it difficult to track and manage. When the SIM is wiped, the international mobile subscriber identity (IMSI) is lost, and afterward the user cannot be recognized. However, you can attempt to blacklist the smartphone through its provider by using the International Mobile Equipment Identity (IMEI), electronic serial number (ESN), or mobile equipment identifier (MEID). The ID used will vary depending on the type and age of smartphone. Regardless, as a security administrator, you should avoid that tactic altogether because the damage has already been done, so protection of the SIM becomes vital.

Skimming

Key Topic

Skimming is a type of attack in which an attacker captures credit card information or information from other similar cards (gift cards, loyalty cards, identification cards, and so on) from a cardholder surreptitiously. Attackers use a device called a *skimmer* that can be installed at strategic locations such as ATMs and gas pumps to collect card data.

Attackers can also perform radio frequency identification (RFID) attacks to steal information. RFID has many uses, but it all boils down to identifying and tracking tags that are attached to objects. In the security world, that generally means authentication. As with any wireless technology, RFID is susceptible to attack. For example, some RFID tags can be affected by skimming, MITM attacks, sniffing, eavesdropping/replaying, spoofing, and jamming (DoS).

From an authentication standpoint, the attacker is using these attacks to try to find out the passcode. An RFID tag can also be reverse-engineered if the attacker gets possession of it. Finally, power levels can be analyzed to find out passwords. On some RFID tags, correct passcodes emit a different level of power than incorrect passcodes. To prevent these attacks, you, as the security administrator, or your team should consider newer-generation RFID devices, encryption, chip coatings, filtering of data, and multifactor authentication methods. Encryption is one of the best

methods. Included in this prevention method are rolling codes, which are generated with a pseudorandom number generator (PRNG) and challenge-response authentication (CRA), where the user (or user's device) must present the valid response to the challenge.

RFID ranges vary depending on the EM band used—from 10 cm up to 200 meters. Many authentication readers can be as much as 1 meter, which is enough to facilitate skimming of information by attackers. One way to avoid this is to use newer RFID proximity readers and keys—ones that use lower frequencies—from respectable companies.

Another way is to utilize the set of protocols called near-field communication (NFC). NFC generally requires that communicating devices be within 4 cm of each other, which makes skimming of information difficult. If an employee uses a smartphone to enter a building instead of an RFID device, then NFC should be implemented. NFC has obvious benefits for contactless payment systems or any other non-contact-oriented communications between devices. However, for optimal security, you should use contact-oriented readers and cards.

Adversarial Artificial Intelligence

Attackers have developed new ways to attack artificial intelligence (AI) and machine learning (ML) implementations. Examples of these adversarial techniques are tainting training data for ML and leveraging insecure implementations of machine learning algorithms.

Tainted Training Data for Machine Learning

Key Topic

Machine learning is being implemented in many solutions in use today (even in security products and solutions). Attackers can manipulate (taint) training data used for machine learning implementations to cause errors in the outcome of the ML solution. They can either change the integrity and modify the training data, or they can inject incorrect data in the training set.

TIP If you are not familiar with the concepts of artificial intelligence and machine learning, the following article provides a good overview: https://vas3k.com/blog/machine_learning.

Security of Machine Learning Algorithms

Key Topic

Securing machine learning algorithms involves protecting the basis of security: confidentiality, integrity, and availability. The CIA concepts that apply to non-ML

implementations also apply here with a few specific techniques. To protect ML implementations, you must classify malicious events in order to prevent them from interfering with algorithms and system operations. You need to prevent attackers from reaching ML system assets and from interfering with normal operations. You also need to prevent the attacker from reading critical data, as well as injecting any malicious or erroneous data in the training set.

Data poisoning attacks can be carried out by attackers to fool the machine learning system to produce incorrect errors. Another attack against machine learning implementations is the overfitting attack. When a machine learning system starts the learning process and "memorizes" its training data set, it will not generalize to new data. This process is referred to as overfitting. Overfitting and overfit models are very easy to attack because adversarial examples need only be a short distance away from the input space in the ML system.

Many ML systems are implemented by tuning a base model that is already trained to enhance the learning process. ML transfer attacks can be launched by an attacker in this scenario. When the pretrained model is widely available, an attacker may be able to formulate attacks using this pretrained model that will be robust enough to succeed against your tuned task-specific model (which is typically unavailable to the attacker). In addition, the ML system you are fine-tuning could possibly be a Trojan injected by the attacker that includes devious ML behavior that is unanticipated.

Supply-Chain Attacks

A ***supply-chain attack*** is a type of attack in which attackers target security weaknesses in the supply network. Attackers have successfully launched supply-chain attacks against many different types of industries (including manufacturing plants, financial services companies, energy and oil companies, technology companies, and more). Attackers can modify products or software during or right after the manufacturing process of a product by installing a rootkit or hardware-based spying components.

Cybersecurity experts recommend strict control of your supply network to prevent potential damage from today's attackers. The supply chain is a complex network of interconnected stakeholders governed by supply and demand.

TIP An example of a supply-chain attack is the manipulation of system maintenance software called CCleaner. This software was manipulated in the supply chain including implants designed to spy and steal sensitive information from many large technology companies. You can obtain more information about this attack at https://blog.talosintelligence.com/2017/09/ccleaner-c2-concern.html.

Cloud-based vs. On-premises Attacks

Many organizations are moving to the cloud or deploying hybrid solutions to host their applications. Organizations moving to the cloud are almost always looking to transition from capital expenditure (CapEx) to operational expenditure (OpEx). Most Fortune 500 companies operate in a multicloud environment. It is obvious that cloud computing security is more important than ever. Cloud computing security includes many of the same functionalities as traditional IT security, which includes protecting critical information from theft, data exfiltration, and deletion, as well as privacy.

NOTE You will learn details about cloud computing architectures in Chapter 10, "Summarizing Virtualization and Cloud Computing Concepts."

Cloud Security Threats

There are many potential threats when organizations move to a cloud model. For example, although your data is in the cloud, it must reside in a physical location somewhere. Your cloud provider should agree in writing to provide the level of security required for your customers. The following are questions to ask a cloud provider before signing a contract for its services:



- **Who has access?** Access control is a key concern because insider attacks are a huge risk. Anyone who has been approved to access the cloud is a potential hacker, so you want to know who has access and how they were screened. Even if it was not done with malice, an employee can leave, and then you find out that you don't have the password, or the cloud service gets canceled because maybe the bill didn't get paid.
- **What are your regulatory requirements?** Organizations operating in the United States, Canada, and the European Union must abide by many regulatory requirements (for example, ISO/IEC 27002, EU-U.S. Privacy Shield Framework, ITIL, and COBIT). You must ensure that your cloud provider can meet these requirements and is willing to undergo certification, accreditation, and review.
- **Do you have the right to audit?** This particular item is no small matter in that the cloud provider should agree in writing to the terms of the audit. With cloud computing, maintaining compliance could become more difficult to achieve and even harder to demonstrate to auditors and assessors. Of the many regulations touching on information technology, few were written with cloud computing in mind. Auditors and assessors might not be familiar with cloud computing generally or with a given cloud service in particular.

Division of compliance responsibilities between cloud provider and cloud customer must be determined before any contracts are signed or service is started.

- **What type of training does the provider offer its employees?** This is a rather important item to consider because people will always be the weakest link in security. Knowing how your provider trains its employees is an important item to review.
- **What type of data classification system does the provider use?** Questions you should be concerned with here include what data classification standard is being used and whether the provider even uses data classification.
- **How is your data separated from other users' data?** Is the data on a shared server or a dedicated system? A dedicated server means that your information is the only thing on the server. With a shared server, the amount of disk space, processing power, bandwidth, and so on is limited because others are sharing this device. If it is shared, the data could potentially become comingled in some way.
- **Is encryption being used?** Encryption should be discussed. Is it being used while the data is at rest and in transit? You will also want to know what type of encryption is being used. For example, there are big technical differences between DES and AES. For both of these algorithms, however, the basic questions are the same: Who maintains control of the encryption keys? Is the data encrypted at rest in the cloud? Is the data encrypted in transit, or is it encrypted at rest and in transit?
- **What are the service-level agreement (SLA) terms?** The SLA serves as a contracted level of guaranteed service between the cloud provider and the customer that specifies what level of services will be provided.
- **What is the long-term viability of the provider?** How long has the cloud provider been in business, and what is its track record? If it goes out of business, what happens to your data? Will your data be returned and, if so, in what format?
- **Will the provider assume liability in the case of a breach?** If a security incident occurs, what support will you receive from the cloud provider? While many providers promote their services as being "unhackable," cloud-based services are an attractive target to hackers.
- **What is the disaster recovery/business continuity plan (DR/BCP)?** Although you might not know the physical location of your services, it is physically located somewhere. All physical locations face threats such as fire, storms, natural disasters, and loss of power. In case of any of these events, how will the cloud provider respond, and what guarantee of continued services is it promising?

Even when you end a contract, you must ask what happens to the information after your contract with the cloud service provider ends.

Insufficient due diligence is one of the biggest issues when moving to the cloud. Security professionals must verify that issues such as encryption, compliance, and incident response are worked out before a contract is signed.

Cloud Computing Attacks

Because cloud-based services are accessible via the Internet, they are open to any number of attacks. As more companies move to cloud computing, look for hackers to follow. Some of the potential attack vectors that criminals might attempt include the following:

Key Topic

- **Session hijacking:** This attack occurs when the attacker can sniff traffic and intercept traffic to take over a legitimate connection to a cloud service.
- **DNS attack:** This form of attack tricks users into visiting a phishing site and giving up valid credentials.
- **Cross-site scripting (XSS):** This type of attack is used to steal cookies that can be exploited to gain access as an authenticated user to a cloud-based service.
- **SQL injection:** This attack exploits vulnerable cloud-based applications that allow attackers to pass SQL commands to a database for execution.
- **Session riding:** This term is often used to describe a cross-site request forgery (CSRF) attack. Attackers use this technique to transmit unauthorized commands by riding an active session using an email or malicious link to trick users while they are currently logged in to a cloud service.
- **Distributed denial-of-service (DDoS) attack:** Some security professionals have argued that the cloud is more vulnerable to DDoS attacks because it is shared by many users and organizations, which also makes any DDoS attack much more damaging.
- **Man-in-the-middle cryptographic attack:** This attack is carried out when the attacker places himself or herself in the communication path between two users. Any time the attacker can do this, there is the possibility that he or she can intercept and modify communications.
- **Side-channel attack:** An attacker could attempt to compromise the cloud by placing a malicious virtual machine in close proximity to a target cloud server and then launching a side-channel attack.

- **Authentication attack:** Authentication is a weak point in hosted and virtual services and is frequently targeted. There are many ways to authenticate users, such as based on what a person knows, has, or is. The mechanisms used to secure the authentication process and the method of authentication used are frequent targets of attackers.
- **API attacks:** Often APIs are configured insecurely. An attacker can take advantage of API misconfigurations to modify, delete, or append data in applications or systems in cloud environments.

Regardless of the model used, cloud security is the responsibility of both the client and the cloud provider. These details will need to be worked out before a cloud computing contract is signed. The contracts will vary depending on the given security requirements of the client. Considerations include disaster recovery, SLAs, data integrity, and encryption. For example, is encryption provided end to end or just at the cloud provider? Also, who manages the encryption keys: the cloud provider or the client? Overall, you need to ensure that the cloud provider has the same layers of security (logical, physical, and administrative) in place that you would have for services you control.

Cryptographic Attacks

Attackers can launch different ***cryptographic attacks*** against weak cryptography (crypto) implementations. The following sections describe three common crypto attacks: collision, birthday, and downgrade attacks.

Collision

Key Topic

A ***collision*** occurs when two different files end up using the same hash. Message Digest Algorithm 5 (MD5) is a legacy hashing algorithm that is used to attempt to provide data integrity. By checking the hash produced by the downloaded file against the original hash, you can verify the file's integrity with a level of certainty. However, MD5 hashes are susceptible to collisions. Due to this low collision resistance, MD5 is considered to be harmful today. MD5 is also vulnerable to threats such as rainbow tables and preimage attacks. The best solution to protect against these attacks is to use a stronger type of hashing function such as SHA-2 or higher. The Secure Hash Algorithm (SHA) is one of a number of hash functions designed by the U.S. National Security Agency (NSA) and published by NIST. These functions are used widely in the U.S. government. Because MD5 and SHA-1 have vulnerabilities, government agencies and the private sector started using SHA-2 and newer implementations.

NOTE You will learn more about hashing algorithms in Chapter 16, "Summarizing the Basics of Cryptographic Concepts."

It is important that a hashing algorithm be collision-resistant. If it has the capability to avoid the same output from two guessed inputs (by an attacker attempting a collision attack), it is collision-resistant. When it comes to cryptography, "perfect hashing" is not possible because usually unknowns are involved, such as the data to be used to create the hash and what hash values have been created in the past. Though perfect is not possible, it is possible to increase collision resistance by using a more powerful hashing algorithm.

Birthday

Key Topic

A ***birthday attack*** is an attack on a hashing system that attempts to send two different messages with the same hash function, causing a collision (similarly to the concept explained in the preceding section). It is based on the birthday problem in probability theory (also known as the birthday paradox). It can be summed up simply as the following: a randomly chosen group of people will have a pair of persons with the same calendar date birthday. Given a standard calendar year of 365 days, the probability of this occurring with 366 people is 100 percent (367 people on a leap year). So far, this makes sense and sounds logical.

The paradox (thoughtfully and mathematically) comes into play when fewer people are involved. With only 57 people, there is a 99 percent probability of a match (a much higher percentage than one would think), and with only 23 people, there is a 50 percent probability. Imagine that and blow out your candles! And by this, I mean use hashing functions with strong collision resistance. Because if attackers can find any two messages that digest the same way (use the same hash value), they can deceive a user into receiving the wrong message. To protect against a birthday attack, you should use a secure transmission medium, such as SSH, or encrypt the entire message that has been hashed.

Downgrade

Key Topic

A ***downgrade attack*** is a type of cryptographic attack that forces the rollback of a strong algorithm in favor of an older, lower-quality algorithm or mode of operation. Attackers leverage systems that have legacy crypto algorithms typically enabled for backward compatibility with older systems. Downgrade attacks can be performed by attackers in combination with an MITM attack.

Downgrade attacks can take many forms. However, they can target the crypto algorithm itself (such as downgrading from AES to DES or RC4) or the version of an algorithm (such as downgrading from TLS 1.3 to 1.0).

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 2-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 2-2 Key Topics for Chapter 2

Key Topic Element	Description	Page Number
Paragraph	Description of malware	33
Paragraph	Description of ransomware	33
Paragraph	Description of cryptomalware	34
Figure 2-1	A ransomware attack	34
Paragraph	Description of Trojan horses	35
Paragraph	Description of remote access Trojans	35
Paragraph	Description of rootkits	35
Paragraph	Description of worms	36
Paragraph	Description of fileless malware	37
Paragraph	Description of Command and Control, bots, and botnets	37
Figure 2-2	A botnet and C2 example	38
Figure 2-3	The botnet performing a DDoS attack	39
Paragraph	Description of logic bombs	39
Paragraph	Description of spyware	40
Paragraph	Description of potentially unwanted programs (PUPs)	40
Paragraph	Description of keyloggers	42
Paragraph	Description of backdoors	43
Paragraph	Description of dictionary-based and brute-force attacks	45
Paragraph	Description of password spraying	46
Paragraph	Description of password cracking	46
Figure 2-4	Offline password cracking	47
Paragraph	Description of rainbow tables	47

Key Topic Element	Description	Page Number
Paragraph	Description of password hacking via pass the hash technique	47
Paragraph	Description of malicious flash drives	48
Paragraph	Description of malicious USB drives	48
Paragraph	Description of card cloning attacks	48
Paragraph	Description of skimming	49
Paragraph	Description of tainted training data used for machine learning	50
Paragraph	Description of securing machine learning algorithms	50
List	Questions to ask a cloud provider before signing a contract for its services	52
List	Attack vectors against cloud-based services	54
Paragraph	Description of collisions and hashing algorithm security vulnerabilities	55
Paragraph	Description of birthday attacks	56
Paragraph	Description of downgrade attacks	56

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

malware, ransomware, cryptomalware, Trojans, remote access Trojans (RATs), rootkit, worm, fileless malware, bots, botnet, command and control, logic bomb, spyware, potentially unwanted programs (PUPs), keylogger, backdoors, dictionary-based attack, brute-force attack, password spraying, online password cracker, offline password cracker, rainbow table, plaintext, malicious flash drives, malicious USB cables, card cloning attacks, skimming, supply-chain attack, cryptographic attacks, collision, birthday attack, downgrade attack

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is a group of compromised computers that have software installed by a worm or Trojan?
2. What term is often used to describe a compromised system that can be updated automatically and remotely?
3. What is a common symptom of spyware?
4. You noticed that your DHCP server is flooded with information. After analyzing this condition, you found that the information is coming from more than 100 computers on the network. What is most likely the reason?
5. Which type of malicious software encrypts sensitive files and asks the user to pay in order to obtain a key recover those files?
6. What is a malicious attack that executes at the same time every week?
7. What is still one of the most common ways that attackers spread ransomware?
8. What is a type of malware that appears to a user as legitimate but actually enables unauthorized access to the user's computer?



This chapter covers the following topics related to Objective 1.3 (Given a scenario, analyze potential indicators associated with application attacks) of the CompTIA Security+ SY0-601 certification exam:

- Privilege escalation
- Cross-site scripting
- Injections
  - Structured query language (SQL)
  - Dynamic link library (DLL)
  - Lightweight directory access protocol (LDAP)
  - Extensible markup language (XML)
- Pointer/object dereference
- Directory traversal
- Buffer overflows
- Race conditions (Time of check/time of use)
- Error handling
- Improper input handling
- Replay attack (session replays)
- Integer overflow
- Request forgeries
  - Server-side
  - Cross-site
- Application programming interface (API) attacks
- Resource exhaustion
- Memory leak
- Secure Socket Layer (SSL) stripping
- Driver manipulation
  - Shimming
  - Refactoring
- Pass the hash

Analyzing Potential Indicators Associated with Application Attacks

Thanks to the advancements in modern web applications and related frameworks, the ways to create, deploy, and maintain web applications have changed such that the environment is now very complex and diverse. These advancements in web applications have also attracted threat actors.

Web-based applications are everywhere. You can find them for online retail, banking, enterprise applications, mobile, and Internet of Things (IoT) applications. In this chapter, you learn about different application-based vulnerabilities and related attacks such as privilege escalation, cross-site scripting, injection vulnerabilities, directory traversal, buffer overflows, race conditions, improper input handling, replay attacks, application programming interface (API) attacks, and many others.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 3-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 3-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Privilege Escalation	1
Cross-Site Scripting (XSS) Attacks	2
Injection Attacks	3
Pointer/Object Dereference	4
Directory Traversal	5

Foundation Topics Section	Questions
Buffer Overflows	6
Race Conditions	7
Error Handling	8
Improper Input Handling	9
Replay Attacks	10–11
Request Forgeries	12
Application Programming Interface (API) Attacks	13
Resource Exhaustion	14
Memory Leaks	15
Secure Socket Layer (SSL) Stripping	16
Driver Manipulation	17
Pass the Hash	18

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What type of privilege escalation attack occurs when a user accesses functions or content reserved for other normal users—for example, if one user reads another's email?
    a. Horizontal
    b. Vertical
    c. Sudo abuse
    d. None of these answers are correct.

2. Which vulnerability can lead to stealing of cookies or redirecting users to malicious sites and where the malicious code or script is permanently stored on a vulnerable system?
    a. DOM-based XSS
    b. Stored XSS
    c. Reflected XSS
    d. All of these answers are correct.

3. Which type of injection attack occurs when code is run within the address space of another process, forcing it to load another library?
   a. DLL injection
   b. LDAP injection
   c. SQL injection
   d. None of these answers are correct.

4. Which condition occurs when a program dereferences a pointer that it expects to be valid but is null, which can cause the application to exit or the system to crash?
   a. Address space layout randomization
   b. Tunnel-trap
   c. Format string
   d. Null pointer dereferences

5. Which type of attack occurs when an attacker tries to escape the web root folder and access arbitrary files by using ../../ in an URL?
   a. Directory traversal
   b. Path traversal
   c. Dot-dot-slash
   d. All of these answers are correct.

6. What occurs when arithmetic operations attempt to create a numeric value that is too big for the available memory space?
   a. Stack overruns
   b. Integer overflows
   c. Format strings
   d. Heap underruns

7. Race conditions are also known as _____ attacks.
   a. Heap overflows
   b. Time-of-check (TOC) or time-of-use (TOU)
   c. Stack overflows
   d. Buffer overflows

8. _____code should be checked thoroughly so that a malicious user can't find out any additional information about the system.
   a. Regression
   b. Patched
   c. Error exception handling
   d. Remote code execution

9. Input validation vulnerabilities can be found by using which of the following techniques?
   a. A python debugger
   b. Fuzzing
   c. Heap flags
   d. All of these answers are correct.

10. Which type of attack occurs when an attacker might use a packet sniffer to intercept data and retransmit it later?
    a. Overflow
    b. DLL injection
    c. XSS
    d. Replay

11. A web session ID is typically represented in a cookie. Which type of attack occurs when an attacker could steal a valid user's session ID and reuse it to perform malicious transactions?
    a. Session replay
    b. Session fuzzing
    c. SQL injection
    d. All of these answers are correct.

12. _____ attacks leverage the trust that the application has in the targeted user. For example, the attacker could inherit the privileges of the user to perform an undesired action, such as stealing sensitive information, creating users, or downloading malware.
    a. XSS
    b. SQLi
    c. XML injection
    d. XSRF

13. Which of the following are general best practices to protect APIs?
    a. Secure API services to provide only HTTPS endpoints with a strong version of TLS.
    b. Validate parameters in the application and sanitize incoming data from API clients.
    c. Explicitly scan for common attack signatures; injection attacks often betray themselves by following common patterns.
    d. All of these answers are correct.

14. Which of the following is a type of denial-of-service (DoS) attack?
    a. C2
    b. Botnet
    c. Resource exhaustion
    d. None of these answers are correct.

15. Which of the following attacks can reduce the performance of a computer, especially in systems with shared memory or limited memory?
    a. Heap underrun
    b. Format string
    c. XSS
    d. Memory leak

16. Which of the following is a common way to launch an SSL strip attack?
    a. Creating a rogue wireless access point (hotspot)
    b. Launching Metasploit
    c. Fuzzing an application
    d. Leveraging a memory leak

17. Which of the following is a driver manipulation attack where an attacker adds a small library that intercepts API calls?
    a. Kernel module attack
    b. Request forgery
    c. Refactoring
    d. Shimming

18. Pass the hash are attacks usually performed in which of the following operating systems?
    a. Apple iOS
    b. macOS
    c. Linux
    d. Windows

Foundation Topics

Privilege Escalation

Key Topic

Privilege escalation is the act of exploiting a bug or design flaw in a software or firmware application to gain access to resources that normally would have been protected from an application or user. This results in a user gaining additional privileges, more than were originally intended by the developer of the application—for example, if a regular user gains administrative control, or if a particular user can read another user's email without authorization.

The original developer does not intend for the attacker to gain higher levels of access but probably doesn't enforce a need-to-know policy properly and/or hasn't validated the code of the application appropriately. This technique is used by attackers to gain access to protected areas of operating systems or to applications—for example, if a particular user can read another user's email without authorization. Buffer overflows are used on Windows computers to elevate privileges as well. To bypass digital rights management (DRM) on games and music, attackers use a method known as jailbreaking, another type of privilege escalation, most commonly found on mobile devices. Malware also attempts to exploit privilege escalation vulnerabilities, if any exist on the system. Privilege escalation can also be attempted on network devices. Generally, the fix for this is simply to update the device and to check on a regular basis if any updates are available. For example, a typical SOHO router has a user account and an administrator account. If a device like this isn't properly updated, an attacker can take advantage of a bug in the firmware to elevate the privileges of the user account. Couple this with the fact that a person forgot to put a password on the user account (or disable it) and your network could be in for some "fun." It is also possible on some devices to encrypt the firmware component. The following are a couple different types of privilege escalation:

Key Topic

- **Vertical privilege escalation:** When a lower-privileged user accesses functions reserved for higher-privileged users—for example, if a standard user can access functions of an administrator. This is also known as privilege elevation and is the most common description. To protect against this situation, you should update the network device firmware. In the case of an operating system, it should again be updated, and use of some type of access control system is also advisable—for example, User Account Control (UAC).
- **Horizontal privilege escalation:** When a normal user accesses functions or content reserved for other normal users—for example, if one user reads another's email. This can be done through hacking or by a person walking

over to other people's computers and simply reading their email! Always have your users lock their computer (or log off) when they are not physically at their desk!

There is also privilege de-escalation, when high-privileged but segregated users can downgrade their access level to access normal users' functions. Sneaky administrators can attempt this to glean confidential information from an organization. It's a two-way street when it comes to security; you should think three-dimensionally when securing your network!

Cross-Site Scripting (XSS) Attacks

Cross-site scripting (commonly known as XSS) vulnerabilities, which have become some of the most common web application vulnerabilities, are achieved using the following attack types:

Key Topic

- **Reflected XSS:** Reflected XSS attacks (nonpersistent XSS) occur when malicious code or scripts are injected by a vulnerable web application using any method that yields a response as part of a valid HTTP request. An example of a reflected XSS attack is a user being persuaded to follow a malicious link to a vulnerable server that injects (reflects) the malicious code back to the user's browser. This causes the browser to execute the code or script. In this case, the vulnerable server is usually a known or trusted site.
- **Stored (persistent) XSS:** Stored, or persistent, XSS attacks occur when the malicious code or script is permanently stored on a vulnerable or malicious server, using a database. These attacks are typically carried out on websites hosting blog posts (comment forms), web forums, and other permanent storage methods. An example of a stored XSS attack is a user requesting the stored information from the vulnerable or malicious server, which causes the injection of the requested malicious script into the victim's browser. In this type of attack, the vulnerable server is usually a known or trusted site.
- **DOM-based XSS:** The Document Object Model (DOM) is a cross-platform and language-independent application programming interface that treats an HTML, XHTML, or XML document as a tree structure. DOM-based attacks are typically reflected XSS attacks that are triggered by sending a link with inputs that are reflected to the web browser. In DOM-based XSS attacks, the payload is never sent to the server. Instead, the payload is only processed by the web client (browser). In a DOM-based XSS attack,

the attacker sends a malicious URL to the victim; after the victim clicks on the link, it may load a malicious website or a site that has a vulnerable DOM route handler. After the vulnerable site is rendered by the browser, the payload executes the attack in the user's context on that site. One of the effects of any type of XSS attack is that the victim typically does not realize that an attack has taken place. DOM-based applications use global variables to manage client-side information. Often developers create unsecured applications that put sensitive information in the DOM (for example, tokens, public profile URLs, private URLs for information access, cross-domain OAuth values, and even user credentials as variables). It is a best practice to avoid storing any sensitive information in the DOM when building web applications.

Successful exploitation of an XSS vulnerability could result in installation or execution of malicious code, account compromise, session cookie hijacking, revelation or modification of local files, or site redirection.

You typically find XSS vulnerabilities in the following:

- Search fields that echo a search string back to the user
- HTTP headers
- Input fields that echo user data
- Error messages that return user-supplied text
- Hidden fields that may include user input data
- Applications (or websites) that display user-supplied data

The following XSS test can be performed from a browser's address bar:

```
javascript:alert("Omar_s_XSS test");
javascript:alert(document.cookie);
```

The following XSS test can be performed in a user input field in a web form:

```
<script>alert("XSS Test")</script>
```

Attackers can use obfuscation techniques in XSS attacks by encoding tags or malicious portions of the script using Unicode so that the link or HTML content is disguised to the end user browsing the site.

TIP The Open Web Application Security Project (OWASP) is a nonprofit organization where numerous security professionals contribute to projects designed to improve the security of web applications, as well as mobile and IoT devices. I strongly suggest leveraging the numerous resources and projects that OWASP provides—not only for the exam but also to further your knowledge about application-based vulnerabilities, mitigations, and how to prevent them. OWASP provides a list of the top 10 vulnerabilities in web applications. You can find this list at https://owasp.org/www-project-top-ten. XSS has been on that list for many years.

OWASP also provides several recommendations and guidelines to help you prevent XSS at https://cheatsheetseries.owasp.org/cheatsheets/Cross_Site_Scripting_Prevention_Cheat_Sheet.html.

Injection Attacks

Attackers can use several types of ***injection attacks*** to manipulate and compromise an application. The following are the most prevalent injection attacks:

- Structured Query Language (SQL) injection attacks
- Dynamic link library (DLL) injection attacks
- Lightweight Directory Access Protocol (LDAP) injection attacks
- Extensible Markup Language (XML) injection attacks

Structured Query Language (SQL) Injection Attacks

Key Topic

SQL injection (SQLi) vulnerabilities can be catastrophic because they can allow an attacker to view, insert, delete, or modify records in a database. In an SQL injection attack, the attacker inserts, or injects, partial or complete SQL queries via the web application. The attacker injects SQL commands into input fields in an application or a URL in order to execute predefined SQL commands.

Let's take a brief look at SQL. As you might know, the following are some of the most common SQL statements (commands):

- **SELECT:** Used to obtain data from a database
- **UPDATE:** Used to update data in a database
- **DELETE:** Used to delete data from a database
- **INSERT INTO:** Used to insert new data into a database
- **CREATE DATABASE:** Used to create a new database

- **ALTER DATABASE:** Used to modify a database
- **CREATE TABLE:** Used to create a new table
- **ALTER TABLE:** Used to modify a table
- **DROP TABLE:** Used to delete a table
- **CREATE INDEX:** Used to create an index or a search key element
- **DROP INDEX:** Used to delete an index

Typically, SQL statements are divided into the following categories:

- Data definition language (DDL) statements
- Data manipulation language (DML) statements
- Transaction control statements
- Session control statements
- System control statements
- Embedded SQL statements

TIP At the W3Schools website, you can find a tool called the Try-SQL Editor that enables you to practice SQL statements in an "online database" (see https://www.w3schools.com/sql/trysql.asp?filename=trysql_select_all). You can use this tool to become familiar with SQL statements and how they may be passed to an application. Another good online resource that explains SQL queries in detail is https://www.geeksforgeeks.org/sql-ddl-dml-tcl-dcl.

Web applications construct SQL statements involving SQL syntax invoked by the application mixed with user-supplied data. The first portion of the SQL statement shown in Figure 3-1 is not shown to the user; typically, the application sends this portion to the database behind the scenes. The second portion of the SQL statement is typically user input in a web form.



FIGURE 3-1 A SQL Query

If an application does not sanitize user input, an attacker can supply crafted input while trying to make the original SQL statement execute further actions in the database. SQL injections can be done using user-supplied strings or numeric input.

TIP I have created a learning environment called WebSploit that you can use to practice some of these attacks in a safe environment. It includes several intentionally vulnerable applications running in Docker containers. Instructions on how to set up WebSploit can be found at https://websploit.org.

Figure 3-2 shows a basic SQL injection attack.

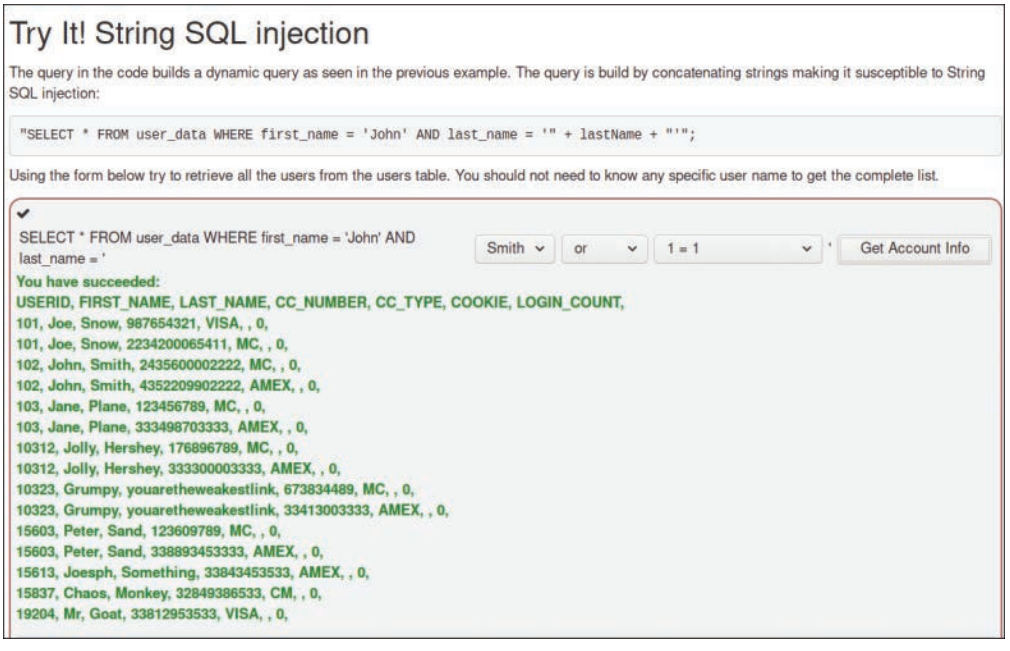


FIGURE 3-2 A Basic SQL Injection Attack

In the figure, WebGoat is used to demonstrate the effects of an SQL injection attack. When the string **Smith' or '1'='1** is entered in the web form, it causes the application to display all records in the database table to the attacker.

One of the first steps when finding SQL injection vulnerabilities is to understand when the application interacts with a database. This is typically done with web authentication forms, search engines, and interactive sites such as e-commerce sites.

You can make a list of all input fields whose values could be used in crafting a valid SQL query. This includes trying to identify and manipulate hidden fields of POST

requests and then testing them separately, plus trying to interfere with the query and to generate an error. As part of penetration testing, you should pay attention to HTTP headers and cookies.

As a penetration tester, you can start by adding a single quote (') or a semicolon (;) to the field or parameter in a web form. The single quote is used in SQL as a string terminator. If the application does not filter it correctly, you might be able to retrieve records or additional information that can help enhance your query or statement.

You can also use comment delimiters (such as -- or /* */), as well as other SQL keywords, including AND and OR operands. Another simple test is to insert a string where a number is expected.

SQL Injection Categories

SQL injection attacks can be divided into the following categories:

Key Topic

- **In-band SQL injection:** With this type of injection, the attacker obtains the data by using the same channel that is used to inject the SQL code. This is the most basic form of an SQL injection attack, where the data is dumped directly in a web application (or web page).
- **Out-of-band SQL injection:** With this type of injection, the attacker retrieves data using a different channel. For example, an email, a text, or an instant message could be sent to the attacker with the results of the query; or the attacker might be able to send the compromised data to another system.
- **Blind (or inferential) SQL injection:** With this type of injection, the attacker does not make the application display or transfer any data; rather, the attacker is able to reconstruct the information by sending specific statements and discerning the behavior of the application and database.

TIP To perform an SQL injection attack, an attacker must craft a syntactically correct SQL statement (query). The attacker may also take advantage of error messages coming back from the application and might be able to reconstruct the logic of the original query to understand how to execute the attack correctly. If the application hides the error details, the attacker might need to reverse-engineer the logic of the original query.

Essentially, five techniques can be used to exploit SQL injection vulnerabilities:

- **Union operator:** This technique is typically used when a SQL injection vulnerability allows a SELECT statement to combine two queries into a single result or a set of results.

- **Boolean:** This technique is used to verify whether certain conditions are true or false.
- **Error-based technique:** This technique is used to force the database to generate an error in order to enhance and refine an attack (injection).
- **Out-of-band technique:** This technique is typically used to obtain records from the database by using a different channel. For example, it is possible to make an HTTP connection to send the results to a different web server or a local machine running a web service.
- **Time delay:** It is possible to use database commands to delay answers. An attacker may use this technique when he or she doesn't get any output or error messages from the application.

It is possible to combine any of the techniques described here to exploit an SQL injection vulnerability. For example, an attacker may use the union operator and out-of-band techniques.

Dynamic Link Library (DLL) Injection Attacks

Key Topic

In ***DLL injection***, code is run within the address space of another process by forcing it to load a dynamic link library (DLL). Ultimately, this type of attack can influence the behavior of a program that was not originally intended. It can be uncovered through penetration testing.

Binary planting is another term used when the attacker implants a malicious binary to a local or remote file system in order for the application to load and execute such malicious binary. You can access detailed information about this technique at OWASP's website at https://owasp.org/www-community/attacks/Binary_planting.

Lightweight Directory Access Protocol (LDAP) Injection Attacks

Key Topic

LDAP injection is similar to SQL injection, again using a web form input box to gain access, or by exploiting weak LDAP lookup configurations. The Lightweight Directory Access Protocol is used to maintain a directory of information such as user accounts or other types of objects. The best way to protect against this type of attack (and all code-injection techniques for that matter) is to incorporate strong input validation.

Extensible Markup Language (XML) Injection Attacks

Key Topic

XML injection attacks can compromise the logic of Extensible Markup Language (XML) applications—for example, XML structures that contain the code for users.

It can be used to create new users and possibly obtain administrative access. You can test for this type of attack by attempting to insert XML metacharacters such as single and double quotes. It can be prevented by filtering *in* allowed characters (for example, A–Z only). This is an example of "default deny" where only what you explicitly filter in is permitted; everything else is forbidden.

One thing to remember is that when attackers utilize code-injecting techniques, they are adding their own code to existing code or are inserting their own code into a form. A variant of this is command injection, which doesn't utilize new code; instead, an attacker executes system-level commands on a vulnerable application or OS. The attacker might enter the command (and other syntax) into an HTML form field or other web-based form to gain access to a web server's password files.

XML External Entity (XXE) is another type of input validation vulnerability and attack against an application that parses XML input. It occurs when input containing a reference to an XML external entity is handled by a poorly configured XML parser. Successful exploitation could lead to the disclosure of confidential data, server-side request forgery, and denial of service. To obtain additional information about this attack and vulnerability, see https://owasp.org/www-community/attacks/.

Pointer/Object Dereference

Key Topic

Another potential memory-related issue deals with pointer dereferencing—for example, the null pointer dereference. Pointer dereferencing is common in programming; when you want to access data (say, an integer) in memory, dereferencing the pointer would retrieve different data from a different section of memory (perhaps a different integer). Programs that contain a ***null pointer dereference*** generate memory fault errors (memory leaks). A null pointer dereference occurs when the program dereferences a pointer that it expects to be valid but is null, which can cause the application to exit or the system to crash. From a programmatical standpoint, the main way to prevent this situation is to use meticulous coding. Programmers can use special memory error analysis tools to enable error detection for a null pointer deference. Once the problem is identified, the programmer can correct the code that may be causing the error(s). But this concept can be used to attack systems over the network by initiating IP address to hostname resolutions—ones that the attacker hopes will fail—causing a return null. What this all means is that the network needs to be protected from attackers attempting this (and many other) programmatic and memory-based attacks via a network connection.

NOTE For more information on null pointer dereferencing (and many other software weaknesses), see the Common Weakness Enumeration portion of MITRE at https://cwe.mitre.org/. Add that site to your favorites!

A programmer can make use of ***address space layout randomization (ASLR)*** to help prevent the exploitation of buffer overflows, remote code execution, and memory corruption vulnerabilities. It randomly arranges the different address spaces used by a program (or process). This can aid in protecting mobile devices (and other systems) from exploits caused by memory-management problems. While there are exploits to ASLR (such as side-channel attacks) that can bypass it and derandomize how the address space is arranged, many systems employ some version of ASLR.

Directory Traversal

Key Topic

Directory traversal, path traversal, or the ../ ("dot-dot-slash") attack, is a method of accessing unauthorized parent (or worse, root) directories. It is often used on web servers that have PHP files and are Linux or UNIX-based, but it can also be perpetrated on Microsoft operating systems (in which case it would be ..\ or the "dot-dot-backslash" attack). It is designed to get access to files such as ones that contain passwords. This access can be prevented by updating the operating system or by checking the code of files for vulnerabilities, otherwise known as fuzzing. For example, a PHP file on a Linux-based web server might have a vulnerable **if** or **include** statement, which when attacked properly could give the attacker access to higher directories and the passwd file.

Figure 3-3 shows an example of a directory traversal (path traversal) attack. The attacker is able to retrieve the contents of the /etc/passwd file of a vulnerable web application. You can also practice this attack using the WebSploit environment.

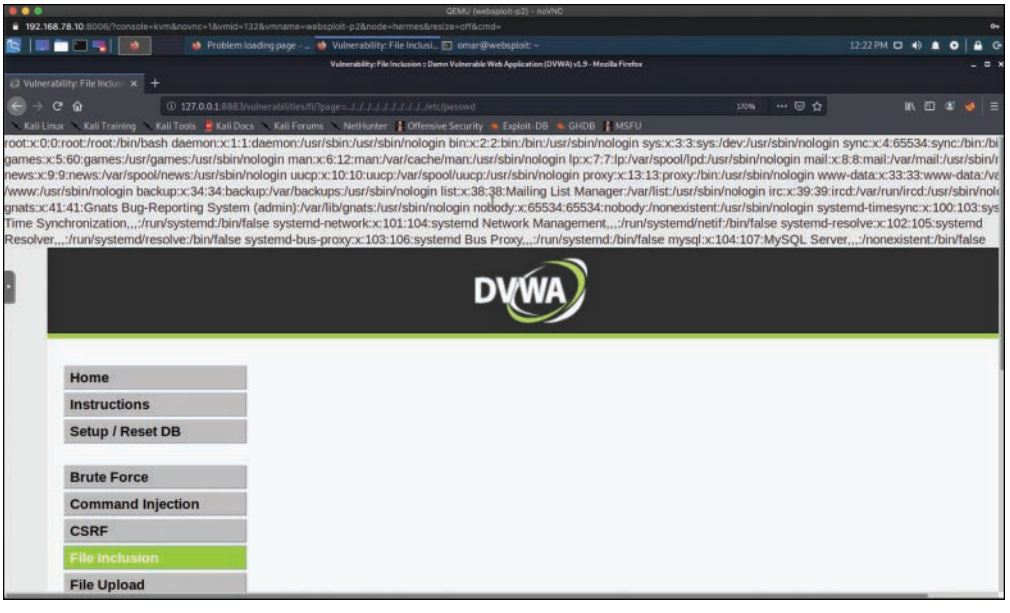


FIGURE 3-3 A Directory Traversal Attack

Buffer Overflows

Key Topic

Memory and buffer vulnerabilities are common. Of the several types, perhaps most important is the buffer overflow. A ***buffer overflow*** occurs when a process stores data outside the memory that the developer intended. This could cause erratic behavior in the application, especially if the memory already had other data in it. Stacks and heaps are data structures that can be affected by buffer overflows. The stack is a key data structure necessary for the exchange of data between procedures. The heap contains data items whose size can be altered during execution. Value types are stored in a stack, whereas reference types are stored in a heap. An ethical coder will try to keep these running efficiently. An unethical coder wanting to create a program vulnerability could, for example, omit input validation, which could allow a buffer overflow to affect heaps and stacks, which in turn could adversely affect the application or the operating system in question.

Let's say a programmer allows for 16 bytes in a string variable. This wouldn't be a problem normally; however, if the programmer failed to verify that *no more than* 16 bytes could be copied over to the variable, that would create a vulnerability that an attacker could exploit with a buffer overflow attack. The buffer overflow can also be initiated by certain inputs. For example, corrupting the stack with no-operation (no-op, NOP, or NOOP) machine instructions, which when used in large numbers can start a NOP slide, can ultimately lead to the execution of unwanted arbitrary code, or can lead to a denial of service on the affected computer.

All this can be prevented by patching the system or application in question, making sure that the OS uses data execution prevention, and utilizing bounds checking, which is a programmatic method of detecting whether a particular variable is within design bounds before it is allowed to be used. It can also be prevented by using correct code, checking code carefully, and using the right programming language for the job in question (the right tool for the right job, yes?). Without getting too much into the programming side of things, special values called "canaries" are used to protect against buffer overflows.

Key Topic

On a semi-related note, ***integer overflows*** occur when arithmetic operations attempt to create a numeric value that is too big for the available memory space. This creates a *wrap* and can cause resets and undefined behavior in programming languages such as C and C++. The security ramification is that the integer overflow can violate the program's default behavior and possibly lead to a buffer overflow. This situation can be prevented or avoided by making overflows trigger an exception condition or by using a model for automatically eliminating integer overflow, such as the CERT As-if Infinitely Ranged (AIR) integer model. You can learn more about this model at https://www.sei.cmu.edu/about/divisions/cert/index.cfm.

Key Topic

Then there are memory leaks. A ***memory leak*** is a type of resource leak caused when a program does not release memory properly. The lack of freed-up memory can reduce the performance of a computer, especially in systems with shared memory or limited memory. A kernel-level leak can lead to serious system stability issues. The memory leak might happen on its own due to poor programming, or it could be that code resides in the application that is vulnerable and is later exploited by an attacker who sends specific packets to the system over the network. This type of error is more common in languages such as C or C++ that have no automatic garbage collection, but it could happen in any programming language. Several memory debuggers can be used to check for leaks. However, it is recommended that garbage collection libraries be added to C, C++, or other programming language, to check for potential memory leaks.

Arbitrary Code Execution/Remote Code Execution

In arbitrary code execution an attacker obtains control of a target computer through some sort of vulnerability, thus gaining the power to execute commands on that remote computer at will. Programs that are designed to exploit software bugs or other vulnerabilities are often called arbitrary code execution exploits. These types of exploits inject "shellcode" to allow the attacker to run arbitrary commands on the remote computer. This type of attack is also known as ***remote code execution (RCE)*** and can potentially allow the attacker to take full control of the remote computer and turn it into a zombie.

RCE commands can be sent to the target computer using the URL of a browser or by using the Netcat service, among other methods. To defend against this type of attack, you should update applications, or if the application is being developed by your organization, it should be checked with fuzz testing and strong input validation (client side and server side) as part of the software development lifecycle (SDLC).

NOTE SDLC is covered in more detail in Chapter 11, "Summarizing Secure Application Development, Deployment, and Automation Concepts."

If you have PHP running on a web server, it can be set to disable remote execution of configurations. A web server (or other server) can also be configured to block access from specific hosts.

NOTE RCE is also very common with web browsers. All browsers have been affected at some point, though some instances are more publicized than others. To see proof of this, access the Internet and search for the Common Vulnerabilities and Exposures (CVE) list for each type of web browser.

Race Conditions

Key Topic

Another exploit is the ***race condition***. This exploit is difficult to perform because it takes advantage of the small window of time between when a service is used and its corresponding security control is executed in an application or OS, or when temporary files are created. It can be defined as anomalous behavior due to a dependence on timing of events. Race conditions are also known as ***time of check (TOC)*** or ***time of use (TOU)*** attacks. Imagine that you are tasked with changing the permissions to a folder or changing the rights in an ACL. If you remove all of the permissions and apply new permissions, then there will be a short period of time when the resource (and system) might be vulnerable. This vulnerability depends on the system used, how it defaults, and how well you have planned your security architecture. That was a basic example, but the race condition is more common within the programming of an application. This exploit can be prevented by proper secure coding of applications and planning of the system and network architecture.

Error Handling

Key Topic

At times, applications will fail. How they fail determines their security. Failure exceptions might show the programming language that was used to build the application, or worse, lead to access holes. ***Error handling*** or error exception handling code should be checked thoroughly so that a malicious user can't find out any additional information about the system. These error handling methods are sometimes referred to technically as pseudocodes. For example, to handle a program exception, properly written pseudocode will basically state (in spoken English): "If a program module crashes, then restart the program module."

Once found, security vulnerabilities should be thoroughly tested, documented, and understood. Patches should be developed to fix the problem but not cause other issues or application regression.

There are some other concepts to consider. First is *obfuscation*, which is the complicating of source code to make it more difficult for people to understand. This is done to conceal its purpose in order to prevent tampering and/or reverse engineering. It is an example of security through obscurity. Other examples of security through obscurity include code camouflaging and steganography, both of which might be used to hide the true code being used. Another important concept is *code checking*, which involves limiting the reuse of code to that which has been approved for use and removing dead code. It's also vital to incorporate good memory management techniques. Finally, you should be very careful when using third-party libraries and software development kits (SDKs) and test them thoroughly before using them within a live application.

For the Security+ exam, the most important of the SDLC phases are maintenance and testing. In the maintenance phase, which doesn't end until the software is removed from all computers, an application needs to be updated accordingly, corrected when it fails, and constantly monitored. It's imperative that you know some of the vulnerabilities and attacks to a system or application, and how to fix them and protect against them. The best way to prevent these attacks is to test and review code.

Key Topic Improper Input Handling

Programmers have various ways to test their code, including system testing, ***input validation***, and fuzzing. When a combination of these testing techniques is used during development and testing, there is a much higher probability of a secure application as the end result.

System testing is generally broken down into two categories: black-box and white-box. ***Black-box testing*** utilizes people who do not know the system. These people (or programs, if automated) test the functionality of the system. Specific knowledge of the system code—and programming knowledge in general—is not required. The tester does not know about the system's internal structure and is often given limited information about what the application or system is supposed to do. In black-box testing, one of the most common goals is to crash the program (often done using fuzzing tools or "fuzzers"). If a user is able to crash the program by entering improper input, the programmer has probably neglected to thoroughly check the error handling code and/or input validation. ***White-box testing*** (also known as transparent testing) is a way of testing the internal workings of the application or system. Testers must have programming knowledge and are given detailed information about the design of the system. They are given login details, production documentation, and source code. System testers might use a combination of fuzzing (covered shortly), data flow testing, and other techniques such as stress testing, penetration testing, and sandboxes.

NOTE A third category that has become popular is gray-box testing, where the tester has internal knowledge of the system from which to create tests but conducts the tests the way a black-box tester would—at the user level.

Stress testing is usually done on real-time operating systems, mission-critical systems, and software, and checks if they have the robustness and availability required by the organization. In a penetration test, the tester mimics what an attacker could do in a vulnerable system.

Compile-Time Errors vs. Runtime Errors

Programmers and developers need to test for potential compile-time errors and runtime errors. Compile time refers to the duration of time during which the statements written in any programming language are checked for errors. ***Compile-time errors*** might include syntax errors in the code and type-checking errors. A programmer can check these without actually "running" the program but instead check it in the compile stage when it is converted into machine code.

A ***runtime error*** is a program error that occurs while the program is running. The term is often used in contrast to other types of program errors, such as syntax errors and compile-time errors. Runtime errors might include running out of memory, invalid memory address access, invalid parameter value, or buffer overflows/dereferencing a null pointer (to name a few), all of which can be discovered only by running the program as a user. Another potential runtime error can occur if there is an attempt to divide by zero. These types of errors result in a software exception. Software and hardware exceptions need to be handled properly. Consequently, ***structured exception handling (SEH)*** is a mechanism used to handle both types of exceptions. It enables the programmer to have complete control over how exceptions are handled and provides support for debugging.

Code issues and errors that occur in either compile time or run time could lead to vulnerabilities in the software. However, it's the runtime environment that you are more interested in from a security perspective, because that more often is where the attacker will attempt to exploit software and websites.

Key Topic

Input validation is very important for website design and application development. ***Input validation***, or data validation, is a process that ensures the correct usage of data: it checks the data that is inputted by users into web forms and other similar web elements. Data not being validated correctly can lead to various security vulnerabilities, including sensitive data exposure and the possibility of data corruption. You can validate data in many ways, from coded data checks and consistency checks to spelling and grammar checks, and so on. Whatever data is being dealt with, it should be checked to make sure it is being entered correctly and won't create or take advantage of a security flaw. If validated properly, bad data and malformed data will be rejected. Input validation should be done both on the client side and, more importantly, on the server side. Let's look at an example next.

If an organization has a web page with a PHP-based contact form, the data entered by the visitor should be checked for errors or maliciously typed input. The following PHP code is contained within a common contact form:

```
else if (!preg_match('/^[A-Za-z0-9.-]+$/', $domain))
 {
 // character not valid in domain part
 $isValid = false;
 }
```

This example is a part of a larger piece of code that checks the entire email address a user enters into a form field. This particular snippet of code checks to make sure the user is not trying to enter a backslash in the domain name portion of the email address. This character is not allowed in email addresses and could be detrimental if used maliciously. In the first line within the brackets, note that it says A-Za-z0-9.-, which tells the system what characters are allowed. Uppercase and lowercase letters, numbers, periods, and dashes are allowed, but other characters such as backslashes, dollar signs, and so on are not allowed. The form's supporting PHP files would interpret those other characters as illegitimate data and would not pass them on through the system. The user would receive an error, which is a part of client-side validation. But the more a PHP form is programmed to check for errors, the more it is possible to have additional security holes. Therefore, server-side validation is even more important. Any data that *is* passed on by the PHP form should be checked at the server as well. In fact, an attacker might not even be using the form in question but instead might be attacking the URL of the web page in some other manner. This type of validation can be checked at the server within the database software or through other means. By the way, the concept of combining client-side and server-side validation also goes for pages that utilize JavaScript.

This is just one basic example, but as mentioned previously, input validation is the key to preventing attacks such as SQL injection and XSS. All form fields should be tested for good input validation code, both on the client side and the server side. By combining the two and checking every access attempt, you develop complete mediation of requests.

NOTE Using input validation is one way to prevent *sensitive data exposure*, which occurs when an application does not adequately protect personally identifiable information (PII). You can also prevent this type of exposure by doing the following: making sure that inputted data is never stored or transmitted in clear text; using strong encryption and securing key generation and storage; using HTTPS for authentication; and using a salt, which is random data used to strengthen hashed passwords.

Replay Attacks

Key Topic

A ***replay attack*** is a network attack in which a valid data transmission is maliciously or fraudulently repeated or delayed. It differs from session hijacking in that the original session is simply intercepted and analyzed for later use. In a replay attack, an attacker might use a packet sniffer to intercept data and retransmit it later. In this way, the attacker can impersonate the entity that originally sent the data.

For example, if customers were to log in to a banking website with their credentials while an attacker was watching, the attacker could possibly sniff out the packets that include the usernames and passwords and then possibly connect with those credentials later on. Of course, if the bank uses Secure Socket Layer (SSL) or Transport Layer Security (TLS) to secure login sessions, then the attacker would have to decrypt the data as well, which could prove more difficult. An organization can defend against this attack in several ways. The first is to use session tokens that are transmitted to people the first time they attempt to connect and identify them subsequently. They are handed out randomly so that attackers cannot guess at token numbers. The second way is to implement timestamping and synchronization, as in a Kerberos environment. A third way would be to use a timestamped ***nonce***, a random number issued by an authentication protocol that can be used only one time. You can also implement CHAP-based authentication protocols to provide protection against replay attacks.

NOTE A replay attack should not be confused with SMTP relay, which occurs when one server forwards email to other email servers.

Key Topic

Session replay attacks occur when an attacker steals a user's valid session ID and reuses that ID to perform malicious transactions and activities with a web application.

Session hijacking is the exploitation of a computer session in an attempt to gain unauthorized access to data, services, or other resources on a computer. A few types of session hijacks can occur:

- **Session theft:** This hijack can be accomplished by making use of packet header manipulation or by stealing a cookie from the client computer, which authenticates the client computer to a server. This is done at the application layer, and the cookies involved are often based off their corresponding web applications (such as WWW sessions). This type of attack can be combated by using encryption and long random numbers for the session key and then regeneration of the session after a successful login. The Challenge Handshake Authentication Protocol (CHAP) can also be employed to require clients to periodically reauthenticate. However, session hijacking can also occur at the network layer—for example, TCP/IP hijacking.

NOTE Details about different authentication protocols are covered in Chapter 24, "Implementing Authentication and Authorization Solutions."

- **TCP/IP hijacking:** This is a common type of session hijacking, due to its popularity among attackers. It occurs when an attacker takes over a TCP session between two computers without the need of a cookie or any other type of host access. Because most communications' authentication occurs only at the beginning of a standard TCP session, an attacker can attempt to gain access to a client computer anytime after the session begins. One way would be to spoof the client computer's IP address, then find out the last packet sequence number sent to the server, and then inject data into the session before the client sends another packet of information to the server. Remember the three-way handshake that occurs at the beginning of a session; this is the only authentication that occurs during the session. A synchronization (SYN) packet is sent by the client to the server, then a SYN/ACK packet is sent by the server to the client, and finally, an acknowledgment (ACK) packet is sent by the client to the server. An attacker can jump in anytime after this process and attempt to steal the session by injecting data into the data stream. This is the more difficult part: the attacker might need to perform a DoS attack on the client to stop it from sending any more packets so that the packet sequence number doesn't increase. In contrast, UDP sessions are easier to hijack because no packet sequence numbers exist. Targets for this type of attack include online games and also DNS queries. To mitigate the risk of TCP/IP hijacking, you should employ encrypted transport protocols such as SSL, IPsec, and SSH.
- **Blind hijacking:** This type of hijacking occurs when an attacker blindly injects data into a data stream without being able to see whether the injection was successful. The attacker could be attempting to create a new administrator account or gain access to one.
- **Clickjacking:** This type of hijacking occurs when a user browsing the web is tricked into clicking something different from what the user thought he or she was clicking. It is usually implemented as a concealed link—embedded code or a script on a website that executes when the user clicks that element. For example, a Flash script, when clicked, could cause the user's webcam to turn on without the user's consent. The user is often redirected to the website from a malicious source. This type of attack can be prevented by updating the user's web browser and using third-party add-ons that watch for clickjacking code or scripts. On the server side, web page frames (such as iframes) must be managed carefully. JavaScript-based snippets can be added and content security policies also can be configured to help manage frames.
- **On-path attack (previously known as man-in-the-middle [MITM] and man-in-the-browser [MITB] attacks):** These attacks intercept all data between a client and a server. It is a type of active interception. If successful,

all communications now go through the MITM attacking computer. The attacking computer can at this point modify the data, insert code, and send it to the receiving computer. This type of eavesdropping is successful only when the attacker can properly impersonate each endpoint. Cryptographic protocols such as Secure Socket Layer (SSL) and Transport Layer Security (TLS) address MITM attacks by using a mutually trusted third-party certification authority (CA). These public key infrastructures (PKIs) should use strong mutual authentication such as secret keys and strong passwords.

On-path attacks also can make use of a Trojan (from a proxy location) that infects a vulnerable web browser and modifies web pages and online transactions, in an attempt to ultimately steal money or data. For example, a user might make an online banking transaction, and the user would see confirmation of the exact transaction, but on the banking side, a different amount might have been actually transferred, with some of it going to a different location altogether. This type of attack can be prevented by updating the web browser, using transaction verification (often third-party), and updating the antimalware on the computer in question.

- ***Watering hole attack*:** As discussed in Chapter 1, a watering hole attack is a targeted attack that occurs when an attacker profiles the websites that the intended victim accesses. The attacker then scans those websites for possible vulnerabilities. If the attacker locates a website that can be compromised, the website is then injected with a JavaScript or other similar code injection that is designed to redirect the user when the user returns to that site (also known as a pivot attack). The user is then redirected to a site with some sort of exploit code…and the rest is, well, history. The purpose is to infect computers in the organization's network, thereby allowing the attacker to gain a foothold in the network for espionage or other reasons. Watering hole attacks are often designed to profile users of specific organizations, and as such, an organization should develop policies to prevent these attacks. This can be done by updating antimalware applications regularly and by using secure virtual browsers that have little connectivity to the rest of the system and the rest of the network. To avoid having a website compromised as part of this attack, you, as administrator, should use proper programming methods and scan the website for malware regularly.

Request Forgeries

Key Topic

Cross-site request forgery (XSRF) is a type of vulnerability in which an attacker lures the targeted user to execute unwanted actions on a web application. Threat-performing XSRF attacks leverage the trust that the application has in the targeted user. For example, the attacker could inherit the privileges of the user to perform an

undesired action, such as stealing sensitive information, creating users, or downloading malware.

A ***server-side request forgery (SSRF)***, unlike a CSRF, is initiated from a web server through a vulnerable web application. With a CSRF attack, the user is tricked into doing something that benefits the attacker. In contrast, an SSRF attack is done for the purpose of compromising information from the web server or enabling other attacks, such as bypassing input validation controls or enabling the attacker to execute further commands. As SSRF attack exploits trust relationships.

Application Programming Interface (API) Attacks

Key Topic

Application programming interfaces (APIs) are used everywhere today. A large number of modern applications use some type of API to allow other systems to interact with the application. Unfortunately, many APIs lack adequate controls and are difficult to monitor. The breadth and complexity of APIs also make it difficult to automate effective security testing. Attackers can launch attacks against APIs to steal, delete, or modify data or to perform a denial-of-service condition.

The following are a few methods or technologies behind modern APIs:

- **Simple Object Access Protocol (SOAP):** This standards-based web services access protocol was originally developed by Microsoft and has been used by numerous legacy applications for many years. SOAP exclusively uses XML to provide API services. XML-based specifications are governed by XML Schema Definition (XSD) documents. SOAP was originally created to replace older solutions such as the Distributed Component Object Model (DCOM) and Common Object Request Broker Architecture (CORBA). You can find the latest SOAP specifications at https://www.w3.org/TR/soap.
- **Representational State Transfer (REST):** This API standard is easier to use than SOAP. It uses JSON instead of XML, and it uses standards such as Swagger and the OpenAPI Specification (https://www.openapis.org) for ease of documentation and to encourage adoption.
- **GraphQL:** GraphQL is a query language for APIs that provides many developer tools. GraphQL is now used for many mobile applications and online dashboards, and many different languages support GraphQL. You can learn more about it at https://graphql.org/code.

SOAP and REST use the HTTP protocol; however, SOAP limits itself to a stricter set of API messaging patterns than REST.

An API often provides a roadmap that describes the underlying implementation of an application. API documentation can provide a great level of detail that can be

very valuable to a security professional, as well to attackers. API documentation can include the following:

- **Swagger (OpenAPI):** Swagger is a modern framework of API documentation and development that is the basis of the OpenAPI Specification (OAS). You can obtain additional information about Swagger at https://swagger.io. The OAS specification is available at https://github.com/OAI/OpenAPI-Specification.
- **Web Services Description Language (WSDL) documents:** WSDL is an XML-based language that is used to document the functionality of a web service. The WSDL specification can be accessed at https://www.w3.org/TR/wsdl20-primer.
- **Web Application Description Language (WADL) documents:** WADL is an XML-based language for describing web applications. The WADL specification can be obtained from https://www.w3.org/Submission/wadl.

The following are general best practices and recommendations for securing APIs:

- Secure API services to only provide HTTPS endpoints with a strong version of TLS.
- Validate parameters in the application and sanitize incoming data from API clients.
- Explicitly scan for common attack signatures; injection attacks often betray themselves by following common patterns.
- Use strong authentication and authorization standards.
- Use reputable and standard libraries to create the APIs.
- Segment API implementation and API security into distinct tiers; doing so frees up the API developer to focus completely on the application domain.
- Identify what data should be publicly available and what is sensitive information.
- If possible, have a security expert do the API code verification.
- Make internal API documentation mandatory.
- Avoid discussing company API development (or any other application development) on public forums.

Resource Exhaustion

Key Topic

Resource exhaustion is an attack against availability that is designed to bring the network, or access to a particular TCP/IP host/server, to its knees by flooding it with useless traffic. Resource exhaustion attacks are a form of denial-of-service

attacks. They can also leverage software vulnerabilities such as memory leaks and file descriptor leaks.

Today, most DoS attacks are launched via botnets (a collection of compromised systems), whereas in the past tools such as the Ping of Death or Teardrop may have been used.

Memory Leaks

Key Topic

A ***memory leak*** is a type of resource leak caused when a program does not release memory properly. The lack of freed-up memory can reduce the performance of a computer, especially in systems with shared memory or limited memory. A kernel-level leak can lead to serious system stability issues. The memory leak might happen on its own due to poor programming, or it could be that code resides in the application that is vulnerable and is later exploited by an attacker who sends specific packets to the system over the network. This type of error is more common in languages such as C or C++ that have no automatic garbage collection, but it could happen in any programming language. Several memory debuggers can be used to check for leaks. However, it is recommended that garbage collection libraries be added to C, C++, or other programming language to check for potential memory leaks.

Secure Socket Layer (SSL) Stripping

Transport Layer Security is the protocol that replaced SSL. However, many people still refer to TLS implementations as "SSL." TLS versions older than 1.2 and all SSL versions are susceptible to SSL stripping attacks. The latest version of TLS (at the time of writing) is version 1.3.

Key Topic

An ***SSL stripping attack*** occurs when an attacker performs a man-in-the-middle attack and can redirect a client to an insecure HTTP connection. The attacker still establishes a secure HTTPS connection between herself or himself and the server, as illustrated in Figure 3-4.

Key Topic

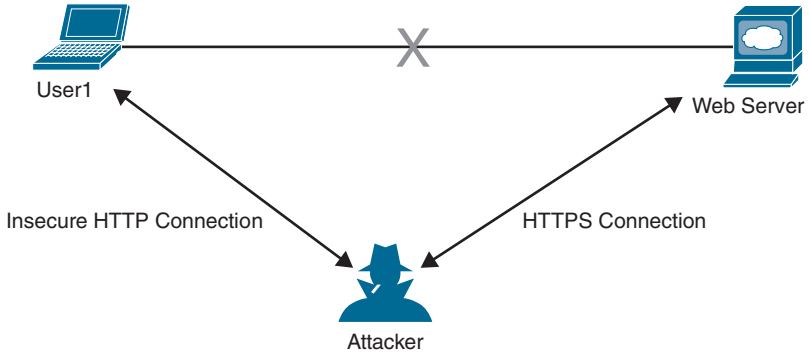


FIGURE 3-4 An SSL Stripping Attack

The attacker in Figure 3-4 downgrades the client connection from HTTPS to HTTP and sends it back to the victim's (User1) browser. The browser is then redirected to the attacker, and all the victim's data will be transferred in plaintext format. Meanwhile, the web server successfully established a secure connection, but with the attacker's machine, not the victim's system.

An attacker can launch an SSL strip attack in many different ways. One of the most common ways is to create a wireless hotspot and lure the victims to connect to it.

Driver Manipulation

Key Topic

It's important to verify that any drivers installed will be compatible with a system. As a security administrator, you should be aware of the potential to modify drivers through the use of driver ***shimming*** (the adding of a small library that intercepts API calls) and driver ***refactoring*** (the restructuring of driver code). By default, shims are not supposed to be used to resolve compatibility issues with device drivers, but it's impossible to foresee the types of malicious code that may present itself in the future. So, you should be careful when updating drivers by making sure that the drivers are signed properly and first testing them on a closed system.

Pass the Hash

All versions of Windows store passwords as hashes in a file called the Security Accounts Manager (SAM) file. The operating system does not know what the actual password is because it stores only a hash of the password. Instead of using a well-known hashing algorithm, Microsoft created its own implementation that has developed over the years.

Microsoft also has a suite of security protocols for authentication, called NT LAN Manager (NTLM). NTLM had two versions: NTLMv1 and NTLMv2. Since Windows 2000, Microsoft has used Kerberos in Windows domains. However, NTLM may still be used when the client is authenticating to a server via IP address or if a client is authenticating to a server in a different Active Directory (AD) forest configured for NTLM trust instead of a transitive interforest trust. In addition, NTLM might also still be used if the client is authenticating to a server that doesn't belong to a domain or if the Kerberos communication is blocked by a firewall.

Key Topic

So, what is a ***pass the hash attack***? Because password hashes cannot be reversed, instead of trying to figure out what the user's password is, an attacker can just use a password hash collected from a compromised system and then use the same hash to log in to another client or server system.

Figure 3-5 shows an example of a pass the hash attack.

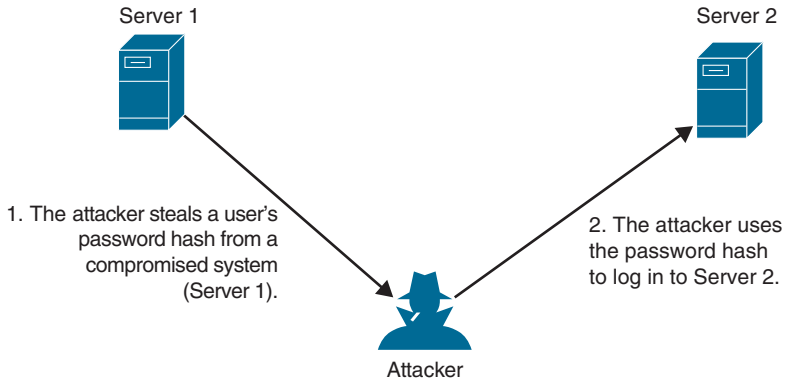


FIGURE 3-5 The Pass the Hash Attack

The Windows operating system and Windows applications ask users to enter their passwords when they log in. The system then converts the passwords into hashes (in most cases using an API called LsaLogonUser). A pass the hash attack goes around this process and just sends the hash to the system to authenticate.

Mimikatz is a tool used by many penetration testers, attackers, and even malware that can be useful for retrieving password hashes from memory; it is a very useful postexploitation tool. You can download the Mimikatz tool from https://github.com/gentilkiwi/mimikatz. Metasploit penetration testing software also includes Mimikatz as a Meterpreter script to facilitate exploitation without the need to upload any files to the disk of the compromised host. You can obtain more information about Mimikatz/Metasploit integration from https://www.offensive-security.com/metasploit-unleashed/mimikatz/.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 3-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 3-2 Key Topics for Chapter 3

Key Topic Element	Description	Page Number
Paragraph	Defining and understanding what privilege escalation is	67
List	Types of privilege escalation attacks	67
List	Understanding the different types of cross-site scripting vulnerabilities	68
Paragraph	Understanding SQL injection	70
List	Identifying SQL injection categories	73
Paragraph	Defining what DLL injection attacks are	74
Paragraph	Understanding LDAP injection attacks	74
Paragraph	Defining XML injection attacks	74
Paragraph	Understanding pointer and object dereference	75
Paragraph	Understanding directory (path) traversal attacks	76
Paragraph	Defining and understanding what buffer overflows are	77
Paragraph	Defining what an integer overflow is	77
Paragraph	Understanding what a memory leak is	78
Paragraph	Defining race conditions	79
Paragraph	Understanding error handling vulnerabilities	79
Section	Improper Input Handling	80
Paragraph	Examples of input validation vulnerabilities	81
Paragraph	Understanding replay attacks	82
Paragraph	Understanding session replay attacks	83
Paragraph	Understanding cross-site request forgery attacks	85
Paragraph	Understanding API-based attacks	86
Paragraph	Defining resource exhaustion attacks	87
Paragraph	Understanding memory leaks	88
Paragraph	Understanding SSL stripping attacks	88
Figure 3-4	An SSL Stripping Attack	88
Paragraph	Understanding driver manipulation, including shimming and refactoring attacks	89
Paragraph	Understanding the pass the hash attack	89
Figure 3-5	The Pass the Hash Attack	90

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

privilege escalation, cross-site scripting, injection attacks, SQL injection (SQLi), DLL injection, LDAP injection, XML injection, null pointer dereference, address space layout randomization (ASLR), directory traversal, buffer overflow, integer overflows, memory leak, remote code execution (RCE), race condition, time of check (TOC) or time of use (TOU), error handling, input validation, black-box testing, white-box testing, compile-time errors, runtime error, structured exception handling (SEH), input validation, replay attack, nonce, session replay, session hijacking, on-path attack, watering hole attack, cross-site request forgery (XSRF), server-side request forgery (SSRF), application programming interfaces (APIs), resource exhaustion, memory leak, SSL stripping attack, shimming, refactoring, pass the hash attack, Mimikatz

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What tool is used by many penetration testers, attackers, and even malware that can be useful for retrieving password hashes from memory?
2. What is the act of restructuring driver code called?
3. What type of attack occurs when the attacker performs an MITM attack and can redirect a client to an insecure HTTP connect?
4. What is a modern framework of API documentation and development that is the basis of the OpenAPI Specification (OAS)?
5. What type of attack occurs when a user browsing the web is tricked into clicking something different than what the user thought he or she was clicking?
6. What type of attack is difficult to exploit because it takes advantage of the small window of time between when a service is used and its corresponding security control is executed in an application, operating system, or when temporary files are created?
7. What feature is supported in most modern operating systems that can help prevent the exploitation of buffer overflows, remote code execution, and memory corruption vulnerabilities?
8. What is a type of input validation vulnerability and attack against an application that parses XML input?

This page intentionally left blank



This chapter covers the following topics related to Objective 1.4 (Given a scenario, analyze potential indicators associated with network attacks) of the CompTIA Security+ SY0-601 certification exam:

- Wireless
  - Evil twin
  - Rogue access point
  - Bluesnarfing
  - Bluejacking
  - Disassociation
  - Jamming
  - Radio frequency identification (RFID)
  - Near-field communication (NFC)
  - Initialization vector (IV)
- On-path attack (previously known as man-in-the-middle attack/ man-in-the-browser attack)
- Layer 2 attacks
  - Address resolution protocol (ARP) poisoning
  - Media access control (MAC) flooding
  - MAC cloning
- Domain name system (DNS)
  - Domain hijacking
  - DNS poisoning
  - Uniform resource locator (URL) redirection
  - Domain reputation
- Distributed denial of service (DDoS)
  - Network
  - Application
  - Operational technology (OT)
- Malicious code or script execution
  - PowerShell
  - Python
  - Bash
  - Macros
  - Visual Basic for Applications (VBA)

Analyzing Potential Indicators Associated with Network Attacks

This chapter covers the potential indicators associated with wired and wireless network attacks. We start with an overview of attacks against wireless networks. Chapter 3, "Analyzing Potential Indicators Associated with Application Attacks," briefly covered on-path attacks (previously known as man-in-the-middle [MITM] or man-in-the-browser [MITB] attacks); however, this chapter covers these topics in more detail. You also learn about the most common attacks against Layer 2 devices and implementation, as well as attacks against DNS. In Chapter 3, you learned about denial-of-service (DoS) and resource exhaustion attacks. In this chapter, you learn what distributed denial-of-service (DDoS) attacks are. The chapter ends with an explanation of how attackers can leverage PowerShell, Python, Bash, macros, and Visual Basic for Applications (VBA) to create malicious code and perform script execution attacks.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 4-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 4-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Wireless Attacks	1
On-Path Attacks	2–3
Layer 2 Attacks	4
Domain Name System (DNS) Attacks	5
Distributed Denial-of-Service (DDoS) Attacks	6
Malicious Code or Script Execution Attacks	7

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is a type of related-key attack, which occurs when an attacker observes the operation of a cipher using several different keys and finds a mathematical relationship between those keys, allowing the attacker to ultimately decipher data?
    a. IV attack
    b. Evil twin
    c. WPA attack
    d. None of these answers are correct.

2. Which of the following can be used to perform an on-path attack?
    a. SQL injection
    b. ARP cache poisoning
    c. Buffer overflow
    d. All of these answers are correct.

3. Which of the following describes when an attacker must first infect the victim's computer with a Trojan or a malicious browser extension or plug-in to intercept transactions from a web browser?
    a. On-path attack
    b. LDAP injection
    c. SQL injection
    d. Keylogger injection

4. Which of the following attacks occurs when the threat actor sends numerous unknown MAC addresses to a network switch to cause a DoS condition or to sniff network connections over a switched network while disrupting network performance?
    a. CAM memory leak attack
    b. MITM
    c. ARP cache poisoning
    d. MAC flooding attacks

5. Which of the following attacks occurs when the attacker changes the registration of a domain name without the permission of the original owner/registrant?
   a. Directory traversal
   b. Path traversal
   c. Dot-dot-slash
   d. Domain hijacking

6. Which of the following protocols has been used for amplification attacks?
   a. IPX
   b. DNS
   c. NetFlow
   d. None of these answers are correct.

7. Which of the following has been used by attackers to create malicious macros in applications such as Excel or Word?
   a. Keyloggers
   b. VBA
   c. DNS
   d. NTP

Foundation Topics

Wireless Attacks

Attackers have performed attacks against wireless networks for many years. These attacks are relevant more than ever because of the advent of Internet of Things (IoT) and mobile devices. In the following sections, you learn the details about these attacks against wireless networks:

- Evil twin
- Rogue access point
- Bluesnarfing
- Bluejacking
- Disassociation
- Jamming
- Radio frequency identifier (RFID)
- Near-field communication (NFC)
- Initialization vector (IV)

Key Topic

Evil Twin Attacks

An ***evil twin*** is a rogue and unauthorized wireless access point (WAP) that uses the same service set identifier (SSID) name as a nearby wireless network, often public hotspots. Like the evil twin antagonist found in many sci-fi books, the device is identical in almost all respects to the authorized WAP. In Figure 4-1, the legitimate (corporate) access point is configured with the SSID corp-net, and the evil twin is deployed by an attacker configured with the same SSID.

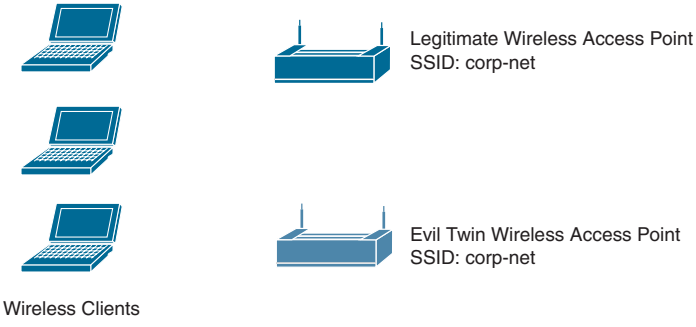


FIGURE 4-1 The Evil Twin Attack

While the antagonist in the sci-fi book usually has a beard or goatee, the WAP is controlled by a person with the same types of motives as the bearded evil twin. One of these motives is to steal sensitive information from wireless users. For example, an attacker might attempt to fool wireless users at an Internet café to connect to the counterfeit WAP to gain valuable information and passwords from the users. If the user is unlucky enough to connect to the evil twin, the attacker can easily record and digest all the information within the session later. This attack can also be enacted on organizations. To protect against this type of attack, organizations can implement virtual private networks (VPNs) that require external authentication outside the WAP. Administrators should scan the network often for rogue APs that might be evil twins. Users in general should be trained not to send passwords, credit card numbers, and other sensitive information over wireless networks.



Rogue Access Points

Rogue APs can be described as unauthorized wireless access points/routers that allow access to secure networks. Sometimes companies lose track of the WAPs on their network. You should keep track of all your devices with network documentation. You also should use network mapping programs and Microsoft Visio to detect and document any rogue APs. Older WAPs, especially ones with weak encryption, should be updated, disabled, or simply disconnected from the network. Some companies may have a dozen WAPs and additional wireless devices such as repeaters, and it may be difficult to keep track of them. In this case, a network mapping program can be one of your best friends. In addition, you can search for rogue APs by using a laptop or handheld computer with the Windows wireless application, the wireless network adapter's built-in software, or third-party applications such as AirMagnet or NetStumbler. If traffic from a rogue AP does enter your network, a network-based intrusion detection system (NIDS) or network-based intrusion prevention system (NIPS) solution (covered in more detail in Chapter 19, "Implementing Secure Network Designs") can be instrumental in detecting and preventing that data and data that comes from other rogue devices. Organizations commonly perform site surveys to detect rogue APs and other unwanted wireless devices.



Bluesnarfing Attacks

Like any wireless technology, Bluetooth is vulnerable to attack. Bluejacking and bluesnarfing are two types of vulnerabilities to Bluetooth-enabled devices. Bluetooth is also vulnerable to conflicts with other wireless technologies. For example, some WLAN (or Wi-Fi) standards use the 2.4-GHz frequency range, as does Bluetooth, and even though Bluetooth uses frequency hopping, conflicts can occur between 802.11g or 802.11b networks and Bluetooth personal area networks (PANs). To avoid this, you should use Bluetooth version 1.2 devices or greater, which employ

adaptive frequency hopping, improving resistance to radio interference. Also, consider placing Bluetooth access points (if they are used) and WLAN access points in different areas of the building. Some companies have policies governing Bluetooth usage; in some cases, it is not allowed if 802.11 standards are in place, and in some cases, a company will enforce rules that say Bluetooth can be used only outside the building. In other cases, a company will put its 802.11 devices on specific channels or use WLAN standards that use the 5-GHz range.

Bluetooth-equipped devices can use near-field communication (NFC), which allows two mobile devices (or a mobile device and a stationary computer) to be automatically paired and transmit data. NFC is not limited to Bluetooth, but Bluetooth is probably the most common technology used to transmit data wirelessly over short distances. Of course, even though the distance is short, it can still be eavesdropped on. In addition, NFC is a data transmission protocol, but not necessarily secure. Data can be destroyed by use of a jammer, and users are also at risk of replay attacks. At the writing of this book, NFC does not offer preventive security in this regard, but a user can prevent these attacks by using only applications that offer SSL/TLS or other secure channels during an NFC session.

I know what you're thinking: the names of these attacks are starting to get a bit ridiculous! I guarantee the naming will improve as we progress through the rest of this book! Anyway, ***bluesnarfing*** is the unauthorized access of information from a wireless device through a Bluetooth connection. Generally, bluesnarfing is the theft of data (calendar information, phonebook contacts, and so on). It is possible to steal other information as well, but to pilfer any of this data, a pairing must be made between the attacking Bluetooth device and the Bluetooth victim. Ways of discouraging bluesnarfing include using a pairing key that is not easy to guess; for example, you should stay away from 0000 or similar default Bluetooth pairing keys! Otherwise, Bluetooth devices should be set to "undiscoverable" (only after legitimate Bluetooth devices have been set up, of course), or Bluetooth can be turned off altogether, especially in areas that might be bluesnarfing playgrounds, such as Times Square in New York City. Some consider bluesnarfing to be a component of bluejacking, but for the Security+ exam, try to differentiate between the two.



Bluejacking Attacks

Bluejacking is the sending of unsolicited messages to Bluetooth-enabled devices such as mobile phones and tablets. Bluejacking is usually harmless, but if it does occur, it may appear that the Bluetooth device is malfunctioning. Originally, bluejackers would send only text messages, but with newer Bluetooth-enabled mobile devices, it is possible to send images and sounds as well. Bluejacking is used in less-than-reputable marketing campaigns. This type of attack can be stopped by setting the affected Bluetooth device to "undiscoverable" or by turning off Bluetooth altogether.

Key Topic

Disassociation and Deauthentication Attacks

DoS attacks can be run against WAPs in a variety of ways. One way is through the use of spoofed MAC addresses. If an attacker emulates a large number of wireless clients, each with a different spoofed MAC address, and never allows authentication for these clients to finish, then legitimate clients will no longer be able to be serviced by the WAP because all of the session spots will have been taken. This is a type of denial-of-service attack due to incomplete authentication. It can be prevented by configuring expiration timeouts for all sessions that have not had activity for a certain period of time. It can also be prevented by updating the WAP and implementing wireless frame protection. (Different providers have different names for this—for example, Cisco Management Frame Protection.)

Another type of wireless-based DoS is the ***Wi-Fi disassociation attack***—also known as Wi-Fi deauthentication attack. The attacker targets a user who is connected to a Wi-Fi network by using a wireless scanner. Then the attacker forces that user's system to deauthenticate using special software and then reauthenticate to the WAP. By capturing the packets during the authentication process, the attacker can find out the SSID or ESSID of the WAP and possibly WPA/WPA2 handshake information. After that, the attacker will attempt to find out the WPA passphrase using a dictionary attack. This process sounds more difficult in theory than it actually is. Using automated software tools, an attacker can accomplish all this in under a minute. These types of "deauth" attacks can be prevented by using WPA2, by using complex passphrases, and by implementing technologies such as management frame protection by Cisco, which adds a hash value to management frames of information, such as the ones used by WPA/WPA2.

WP3 addresses the aforementioned issues using Simultaneous Authentication of Equals (SAE); however, no protocol is perfect. WPA3 also introduces other security vulnerabilities and concerns. For example, the WPA3 Dragonfly handshake used in different Wi-Fi networks is susceptible to different attacks. An attacker can use similar techniques against the EAP implementation and then use the WPA3 Dragonfly handshake to recover a user's password. You can obtain details about these attacks at https://wpa3.mathyvanhoef.com.

A WAP can fall victim to brute-force attacks if WPS is used. The key of the WAP can also be compromised by a brute-force attack known as an exhaustive key search. This type of attack can be prevented by limiting the number of times a password/passphrase can be tried, using time delays between attempts, and requiring complex answers during the authentication process. Also, as a corrective security control, attacking IP addresses can be blocked.

Jamming Attacks

The purpose of ***jamming*** wireless signals or causing wireless network interference is to create a full or partial DoS condition in the wireless network. Such a condition, if successful, is very disruptive. Most modern wireless implementations provide built-in features that can help immediately detect such attacks. To jam a Wi-Fi signal or any other type of radio communication, an attacker basically generates random noise on the frequencies that wireless networks use. With the appropriate tools and wireless adapters that support packet injection, the attacker can cause legitimate clients to disconnect from wireless infrastructure devices.

Radio Frequency Identifier (RFID) Attacks

Radio frequency identification (RFID) has many uses, but it all boils down to identifying and tracking tags that are attached to objects. In the security world, that generally means authentication.

As with any wireless technology, RFID is susceptible to attack. For example, some RFID tags can be affected by skimming, on-path attacks, sniffing, eavesdropping/replaying, spoofing, and jamming (DoS). From an authentication standpoint, the attacker uses these attacks to try to find out the passcode. An RFID tag can also be reverse-engineered if the attacker gets possession of it. Finally, power levels can be analyzed to find out passwords. On some RFID tags, correct passcodes emit a different level of power than incorrect passcodes. To prevent these attacks, a security administrator (or team) should consider newer-generation RFID devices, encryption, chip coatings, filtering of data, and multifactor authentication methods. Encryption is one of the best methods. Included in this prevention method are rolling codes, which are generated with a pseudorandom number generator (PRNG), and challenge-response authentication (CRA), where the user (or user's device) must present the valid response to the challenge.

NOTE As covered in Chapter 2, RFID ranges vary depending on the electromagnetic (EM) frequency band used—from 10 cm up to 200 meters. Many authentication readers can be as much as 1 meter, which is enough to facilitate skimming of information by attackers. One way to avoid this is to use newer RFID proximity readers and keys (ones that use lower frequencies) from respectable companies.

Near-Field Communication (NFC) Attacks

An alternative to RFID is to utilize the set of protocols called ***near-field communication (NFC)***. NFC generally requires that communicating devices be within 4 cm of each other, which makes skimming of information difficult. If an employee

uses a smartphone to enter a building instead of an RFID device, then NFC should be implemented. NFC has obvious benefits for contactless payment systems or any other non-contact-oriented communications between devices. However, for optimal security, you should use contact-oriented readers and cards.

Key Topic

Initialization Vector (IV) Attacks

Initialization vector (IV) attacks are another vulnerability of wireless networks. An IV attack is a type of related-key attack, which occurs when an attacker observes the operation of a cipher using several different keys and finds a mathematical relationship between those keys, allowing the attacker to ultimately decipher data. An initialization vector is a random fixed-sized input that occurs in the beginning of every WEP or WPA packet. For WEP, the IV size was small (24-bit) and led to many successful attacks on WEP. This is why WEP is considered to be deprecated and insecure. The best way to prevent IV attacks is to use stronger wireless protocols such as WPA2 with AES and WPA3.

WEP is susceptible to many different attacks, and it is therefore considered an obsolete wireless protocol. WEP must be avoided, and many wireless network devices no longer support it. WEP keys exist in two sizes: 40-bit (5-byte) and 104-bit (13-byte) keys. In addition, WEP uses a 24-bit IV, which is prepended to the preshared key (PSK). When you configure a wireless infrastructure device with WEP, the IVs are sent in plaintext.

WEP has been defeated for decades. WEP uses RC4 in a manner that allows an attacker to crack the PSK with little effort. The problem is how WEP uses the IVs in each packet. When WEP uses RC4 to encrypt a packet, it prepends the IV to the secret key before including the key in RC4. Consequently, an attacker has the first 3 bytes of an allegedly "secret" key used on every packet. To recover the PSK, the attacker just needs to collect enough data from the air. The attacker can accelerate this type of attack by just injecting ARP packets (because the length is predictable), which allows him or her to recover the PSK much faster. After recovering the WEP key, the attacker can use it to access the wireless network. The attacker can also use the Aircrack-ng set of tools to crack (recover) the WEP PSK.

On-Path Attacks

In an ***on-path attack***, an attacker places himself or herself in-line between two devices or individuals who are communicating in order to eavesdrop or manipulate the data being transferred. As discussed earlier in this chapter, on-path attacks were previously known as man-in-the-middle (MITM) attacks.

On-path attacks can happen at Layer 2 or Layer 3 of the Open Systems Interconnection (OSI) model for computer networks.

TIP If you are not familiar with the OSI model, the following article and embedded video describe it in detail: https://www.networkworld.com/article/3239677/the-osi-model-explained-and-how-to-easily-remember-its-7-layers.html.

Figure 4-2 demonstrates an on-path attack.

Key Topic

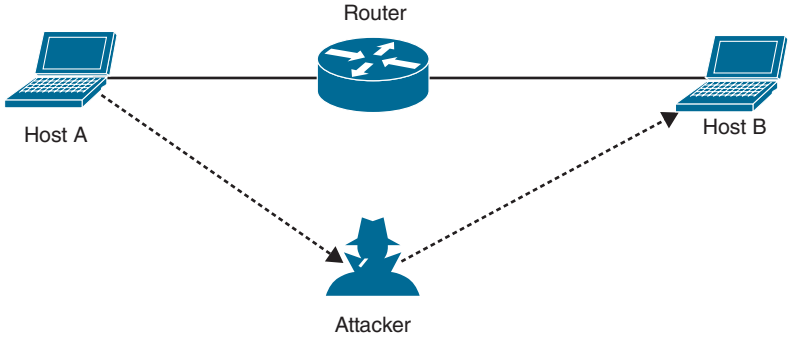


FIGURE 4-2 An On-Path Attack

In the figure the attacker intercepts the packets between host A and host B. On-path attacks could be used to sniff any unencrypted sensitive data or to interact with the user to potentially provide sensitive information to the attacker. ARP cache poisoning attacks can be used to perform on-path attacks. You learn about ARP cache poisoning attacks later in this chapter.

On-path attacks are also often leveraged by attackers in wireless networks. For instance, an attacker could impersonate a wireless access point or "provide free Internet access" in a coffee shop, hotel, airport, or any other public area. Then the attacker could steal information from the users connecting to such a fake wireless access point.

An attacker could also infect a victim's computer with a Trojan or a malicious browser extension or plug-in. The attacker usually gets the malware onto the victim's computer through some form of trickery or deceit. For example, the victim may have been asked to install some plug-in to watch a video or maybe update a program or install a screensaver. After the victim is tricked into installing malware onto his or her system, the malware simply waits for the victim to visit a targeted site. This browser-based on-path attack malware can invisibly modify transaction information like the amount or destination. It can also create additional transactions without the user knowing. Because the requests are initiated from the victim's computer, it is very difficult for the web service to detect that the requests are fake.

Layer 2 Attacks

Layer 2 switched environments (physical or virtual environments) can be targets for attackers. The following sections cover some of the most common Layer 2 attacks.

Address Resolution Protocol (ARP) Poisoning Attacks

ARP cache poisoning (also known as ARP spoofing) is an example of an attack that leads to an on-path attack scenario. An ARP spoofing attack can target hosts, switches, and routers connected to a Layer 2 network by poisoning the ARP caches of systems connected to the subnet and by intercepting traffic intended for other hosts on the subnet. In Figure 4-3, the attacker spoofs Layer 2 MAC addresses to make the victim (the host with the IP address 10.1.1.2) believe that the Layer 2 address of the attacker is the Layer 2 address of its default gateway (10.1.1.1).

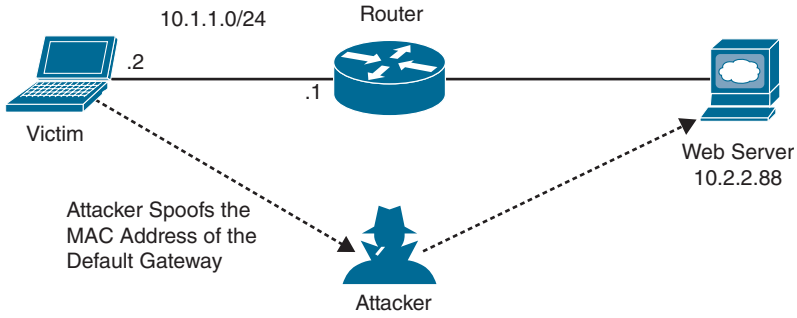


FIGURE 4-3 ARP Cache Poisoning Attack

The packets that are supposed to go to the default gateway are forwarded by the switch to the Layer 2 address of the attacker on the same network. The attacker can forward the IP packets to the correct destination to allow the client to access the web server (10.2.2.88).

A common mitigation for ARP cache poisoning attacks is to use Dynamic ARP Inspection (DAI) on switches to prevent spoofing of the Layer 2 addresses.

Another example of a Layer 2 on-path attack is to place a rogue switch in the network and use Spanning Tree Protocol (STP) to become the root switch. This could allow an attacker to see any traffic that needs to be sent through the root switch.

An attacker can carry out an on-path attack at Layer 3 by placing a rogue router on the network and then tricking the other routers into believing that this new router has a better path. It is also possible to perform an on-path attack by compromising the victim's system and installing malware that can intercept the packets sent by the victim. The malware can capture packets before they are encrypted if the victim is using SSL/TLS/HTTPS or any other mechanism.

Media Access Control (MAC) Flooding Attacks

Key Topic

Layer 2 switches flood frames destined to unknown destination MAC addresses and store the source address and port of every arriving packet in content addressable memory (CAM). All switches have an absolute amount of memory space for the number of MAC addresses that can be learned. Some can scale to millions of entries, whereas others do not. The forwarding table, however, has only a limited address space. Attackers can try to flood the forwarding table to cause a denial-of-service condition. In short, ***MAC flooding attacks*** send numerous unknown MAC addresses to a network switch to cause a DoS condition. In addition, once the Layer 2 forwarding table limit is exceeded, packets are flooded to all ports in a virtual LAN (VLAN). This, in turn, enables the attacker to sniff network connections over a switched network while disrupting network performance.

TIP Port security is a dynamic feature that can be used to limit and identify the MAC addresses of hosts in a network and protect against this type of attack. When a network switch port is configured with port security and the maximum number of MAC addresses is reached, the MAC address of a host attempting to access a port that is different from any of the identified "secure MAC addresses" triggers a security alert to the administrator.

MAC Cloning Attacks

Key Topic

A ***MAC cloning attack*** (or MAC spoofing attack) occurs when an attacker sniffs the network for valid MAC addresses and then uses those MAC addresses to perform other actions. These actions could be to potentially bypass Layer 2 access control policies or masquerade as the default gateway to perform a on-path attack.

Best Practices to Protect Against Layer 2 Attacks

The following are some Layer 2 security best practices for securing your infrastructure:

- Select an unused VLAN (other than VLAN 1) and use it as the native VLAN for all your trunks. Do not use this native VLAN for any of your enabled access ports. Avoid using VLAN 1 anywhere because it is the default.
- Administratively configure switch ports as access ports so that users cannot negotiate a trunk; also disable the negotiation of trunking (that is, do not allow Dynamic Trunking Protocol [DTP]).
- Limit the number of MAC addresses learned on a given port with the port security feature.

- Control Spanning Tree to stop users or unknown devices from manipulating it. You can do so by using the BPDU Guard and Root Guard features.
- Turn off Cisco Discovery Protocol (CDP) on ports facing untrusted or unknown networks that do not require CDP for anything positive. (CDP operates at Layer 2 and might provide attackers information you would rather not disclose.)
- On a new switch, shut down all ports and assign them to a VLAN that is not used for anything other than a parking lot. Then bring up the ports and assign correct VLANs as the ports are allocated and needed.
- Use Root Guard to control which ports are not allowed to become root ports to remote switches.
- Use Dynamic ARP Inspection (DAI).
- Use IP Source Guard to prevent spoofing of Layer 3 information by hosts.
- Implement 802.1X when possible to authenticate and authorize users before allowing them to communicate to the rest of the network.
- Use DHCP snooping to prevent rogue DHCP servers from impacting the network.
- Use storm control to limit the amount of broadcast or multicast traffic flowing through a switch.
- Deploy access control lists, such as Layer 3 and Layer 2 ACLs, for traffic control and policy enforcement.

Domain Name System (DNS) Attacks

DNS is used practically everywhere. For example, every time that you visit a website, the domain name (for example, facebook.com, twitter.com, yourbank.com, and so on) is "resolved" to an IP address. In many cases, when machines communicate to each other using APIs, DNS is also used. Consequently, it is a big target for attackers. In addition to attacks against a DNS infrastructure, attackers have used DNS tunnels to steal data and registered domains to perform malicious tasks (such as exfiltrating data to a malicious domain, hosting malware, and allowing compromised devices to communicate with a command-and-control [C2] server). Let's start by examining how attackers can hijack a domain.

Domain Hijacking Attacks

Domain hijacking is a type of hijacking attack in which the attacker changes the registration of a domain name without the permission of the original owner/registrant. One of the most common methods to perform domain hijacking is using social engineering. An attacker may pretend to be an employee of an organization or the domain registrar over a phone conversation to obtain access to the domain administration area. The attacker could also infect your system with malware (e.g., a keylogger or a Trojan) and steal the credentials for the domain admin.

To prevent domain hijacking, individuals and organizations should use trusted domain registrars. If the registrar provides you with the option to use multifactor authentication, use it! Many individuals and organizations select domain registrars that allow you to anonymize the administrative contact information about a domain (displayed in the WHOIS information).

NOTE WHOIS is a free and publicly accessible directory from which domain names can be queried to discover contact and technical information behind registered domain names. You can access it at https://who.is/.

DNS Poisoning Attacks

DNS poisoning (or DNS cache poisoning) is the modification of name resolution information that should be in a DNS server's cache. It is done to redirect client computers to incorrect websites. This can happen through improper software design, misconfiguration of name servers, and maliciously designed scenarios exploiting the traditionally open architecture of the DNS system. Let's say a client wants to go to www.comptia.org. That client's DNS server will have a cache of information about domain names and their corresponding IP addresses. If CompTIA's site were visited in the recent past by any client accessing the DNS server, its domain name and IP should be in the DNS server's cache. If the cache is poisoned, it could be modified in such a way to redirect requests for www.comptia.org to a different IP address and website. This other site could be a phishing site or could be malicious in some other way. This attack can be countered by using Transport Layer Security (TLS) and digital signatures or by using Domain Name System Security Extensions (DNSSEC), which uses encrypted electronic signatures when passing DNS information, and finally, by patching the DNS server. You might use a Transaction Signature (TSIG) to provide authentication for DNS database updates. This protocol uses shared secret keys and one-way hashing to provide security. One item of note: the hashing procedure might not be secure enough for your organization, so you may want to consider alternatives when updating DNS databases.

Unauthorized zone transfers are another bane to DNS servers. Zone transfers replicate the database that contains DNS data; they operate on top of TCP. If a zone transfer is initiated, say through a reconnaissance attack, server name and IP address information can be stolen, resulting in the attacker accessing various hosts by IP address. To defend against this, zone transfers should be restricted and audited in an attempt to eliminate unauthorized zone transfers and to identify anyone who tries to exploit the DNS server in this manner. Vigilant logging of the DNS server and the regular checking of DNS records can help detect unauthorized zone transfers.

A Windows computer's hosts file can also be the victim of attack. The hosts file is used on a local computer to translate or resolve hostnames to IP addresses. This is the predecessor to DNS, and although the file is normally empty of entries, it is still read and parsed by Windows operating systems. Attackers may attempt to hijack the hosts file in an attempt to alter or poison it or to try to have the client bypass DNS altogether. The best defense for this is to modify the computer's hosts file permissions to read-only. It is located at the following path: \%systemroot%\System32\drivers\etc.

If the file has already been hijacked, and you don't use the file for any static entries, delete it, and Windows should re-create it automatically at the next system startup. If Windows does not, then you can easily re-create a standard hosts file by simply making a blank hosts.txt file and placing it in the path mentioned previously. Some people use the hosts file as a security measure as well. This is done by adding entries that redirect known bad domains to other safe locations or the localhost. Generally, this is done in conjunction with disabling the DNS client service. However, in general, the average Windows user requires the DNS client service.

Hosts files and vulnerable DNS software can also be victims of pharming attacks. As discussed in Chapter 1, ***pharming*** occurs when an attacker redirects one website's traffic to another website that is bogus and possibly malicious. Pharming can be prevented by carefully monitoring DNS configurations and hosts files. Unfortunately, if an ISP's DNS server is compromised, that will be passed on to all the small office/home office (SOHO) routers that the ISP services. So, it becomes more important for end users to be conscious of pharming. They can prevent it by turning on phishing and pharming filters within the browser and by being careful of which websites they access. Users can also check their local hosts files. By default, the file doesn't have any entries, so if they see entries and have never modified the file themselves, they should either delete the entries or delete the file entirely.

Although it is less of an actual attack, ***domain name kiting*** (or simply domain kiting) is the process of deleting a domain name during the five-day grace period (known as the add grace period, or AGP) and immediately reregistering it for another five-day period. This process is repeated any number of times with the end result of having

the domain registered without ever actually paying for it. It is a malicious attack on the entire Domain Name System that misuses the domain-tasting grace period. The result is that a legitimate company or organization often cannot secure the domain name of its choice.

As you can see, the DNS server can be the victim of many attacks due to its visibility on the Internet. It should be closely monitored at all times. Other highly visible servers such as web servers and mail servers should be likewise monitored, audited, and patched as soon as updates are available.

Uniform Resource Locator (URL) Redirection Attacks

Key Topic

Attackers can use social engineering and other techniques to redirect users from a trusted website to a malicious site (***URL redirection attacks***). These techniques include phishing, spear phishing, pharming, and others.

In Chapter 3, you learned about cross-site scripting (XSS) and how attackers can leverage those vulnerabilities to redirect a user to a malicious site. XSS is a flaw in the web application or a website, but the victim is the actual user of such application. These types of attacks are typically the entry point for an attacker into an organization or a user's sensitive data. By validating the input of URLs passed and using whitelists, you can mitigate URL redirection.

Domain Reputation

Key Topic

Domain reputation is a technique used to validate the authenticity of a domain and the services using such domain (including websites and email messages). Domain reputation is used to track known malicious domains that point to websites hosting malware or those that are used for spam, phishing, spear phishing, and other malicious activities. Several technologies and standards have been created for domain authentication and validation. The first example we cover here is the ***Domain Keys Identified Mail (DKIM)***. DKIM is an industry standard defined in RFC 5585, and it provides a means for gateway-based cryptographic signing of outgoing messages. This allows you to embed verification data in an email header and for email recipients to verify the integrity of the email messages.

DKIM uses DNS TXT records to publish public keys. A few additional specifications related to DKIM exist. RFC 6376 introduces DKIM signatures; RFC 5863 provides information on DKIM development, deployment, and operation; and RFC 5617 addresses Author Domain Signing Practices (ADSP).

Another technology is the Sender Policy Framework (SPF), which enables recipients to verify the sender's IP addresses by looking up DNS records that list authorized mail gateways for a particular domain. SPF is an industry standard defined in RFC 4408. SPF uses DNS TXT resource records. Commercial products such as the Cisco Email Security Appliance (ESA) support SPF to verify HELO/EHLO and MAIL FROM identity (FQDN). When you enable SPF, the appliance adds headers in the message, allowing you to obtain additional intelligence on the email sender. One challenge is that the effectiveness of SPF implementations is based on participation. Also, you need to invest time to make sure SPF records are up to date. Many organizations implement SPF because some of the most prevalent email providers nowadays do not accept mail without SPF records.

Some organizations go to the extent of not allowing any emails that do not have an SPF record. Doing so will block more spam emails; however, some legitimate mail might also be dropped if the sending entity hasn't configured SPF correctly.

Then there is also ***Domain-based Message Authentication, Reporting & Conformance (DMARC)***. DMARC is a standard that was designed to thwart spammers from spoofing your domain to send email. Spammers can counterfeit the "From" address on an email message so that it appears to come from a user in your domain. For example, attackers can spoof your bank's domain name (that is, accounts@yourbank.com) to send you a fraudulent email in an effort to obtain your account information. DMARC is designed to block these emails before they appear in your email inbox. DMARC also provides visibility and reports into who is sending email on behalf of your domain to make sure that only legitimate emails are received.

DMARC is an open and free technology, allowing organizations and individuals to secure their domain's emails and maintain control of email delivery.

Distributed Denial-of-Service (DDoS) Attacks

Key Topic

A ***distributed denial-of-service (DDoS) attack*** occurs when a group of compromised systems attacks a single target, causing a DoS to occur at that host. A DDoS attack often utilizes a ***botnet***, which is a large group of computers known as robots or simply "bots." Often, these are systems owned by unsuspecting users. The computers in the botnet that act as attackers are known as zombies. An attacker starts the DDoS attack by exploiting a single vulnerability in a computer system and making that computer the zombie master, or DDoS master. The master system communicates with the other systems in the botnet. The attacker often loads malicious software on many computers (zombies). The attacker then can launch a flood of attacks by all

zombies in the botnet with a single command. DDoS attacks and botnets are often associated with exploit kits (such as the Blackhole kit) and ransomware.

DoS and DDoS attacks are difficult to defend against. Other than the methods mentioned previously in the DoS section, these attacks can be prevented to some extent by updated stateful firewalls, switches, and routers with access control lists; intrusion prevention systems (IPSs); and proactive testing. Several companies offer products that simulate DoS and DDoS attacks. By creating a test server and assessing its vulnerabilities with simulated DoS tests, you can find holes in the security of your server before you take it live. A quick web search for "DoS testing" shows a few of these simulation test companies.

An organization could also opt for a "clean pipe," which attempts to weed out DDoS attacks, among other attacks. This solution is offered as a service by Verisign and other companies. Manual protection of servers can be a difficult task; to implement proper DDoS mitigation, your organization might want to consider anti-DDoS technology and emergency response from an outside source or from the organization's cloud-based provider. Finally, if you do realize that a DDoS attack is being carried out on your network, call your ISP and request that this traffic be redirected.

One specific type of DDoS is the ***DNS amplification attack***. Amplification attacks generate a high volume of packets ultimately intended to flood a target website. In the case of a DNS amplification attack, the attacker initiates DNS requests with a spoofed source IP address. The attacker relies on *reflection*; responses are not sent back to the attacker but are instead sent "back" to the victim server. Because the DNS response is larger than the DNS request (usually), it *amplifies* the amount of data being passed to the victim. An attacker can use a small number of systems with little bandwidth to create a sizable attack. However, a DNS amplification attack can also be accomplished with the aid of a botnet, which has proven to be devastating to sections of the Internet during the period when the attack was carried out.

NOTE Other protocols, such as the Network Time Protocol (NTP), have also been used to perform amplification attacks.

The primary way of preventing this attack is to block spoofed source packets. It can also be prevented by blocking specific DNS servers, blocking open recursive relay servers, rate limiting, and updating your own DNS server(s) often. Finally, as previously mentioned, you can make use of the DNSSEC, which are specifications that provide for origin authentication and data integrity.

Key Topic

Network DDoS attacks aim to target network infrastructure resources (for example, bandwidth, CPU, and memory utilization of the underlying network infrastructure). ***Application DDoS attacks*** target the resources of Layer 7 applications and often leverage known vulnerabilities against specific software. For example, an attacker can exploit denial-of-service vulnerabilities in web servers like Apache HTTP Server ("httpd") and NGINX or a specific web-based application.

Key Topic

Operational technology (OT) is the term used to describe physical items that can be programmed and connected to a network or the Internet. Typically, these devices are used to control electrical grids, pipelines, automobiles, manufacturing plant robots, and other critical infrastructure. Threat actors are always trying to find ways to launch denial-of-service attacks against OT environments and critical infrastructure.

At the same time, Internet of Things devices such as video doorbells, smart thermostats, IP-enabled cameras, and others have been compromised and used to launch very effective DDoS attacks against numerous organizations. A good example is the Mirai botnet that was used to launch a major DDoS several years ago against several high-profile companies and organizations.

Key Topic

Malicious Code or Script Execution Attacks

Attackers have used scripting languages such as ***PowerShell***, ***Python***, and ***Bash*** to perform enumeration, create exploits, and also employ postexploitation techniques. Numerous pieces of malware have used these legitimate tools, as well as ***macros*** embedded in Microsoft Word or Excel documents. A malicious macro is usually placed in documents and emailed to users in the hopes that the users will open the documents, thus executing it.

Visual Basic for Applications (VBA) is an event-driven programming capability in Microsoft operating systems and applications. When you write code using VBA, such code is compiled to a Microsoft proprietary pseudo-code. Applications such as Excel, Word, PowerPoint, and Outlook store this code as a separate stream in COM Structured Storage files—for example, in files such as .xls, .doc, .docx, and .pptx. Attackers have used VBA to create malicious macros and embed them in Excel, Word, and PowerPoint documents. These malicious macros have been designed to steal sensitive information, install keyloggers and Trojans, and perform other nefarious activities.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 4-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 4-2 Key Topics for Chapter 4

Key Topic Element	Description	Page Number
Section	Evil Twin Attacks	98
Section	Rogue Access Points	99
Section	Bluesnarfing Attacks	99
Section	Bluejacking Attacks	100
Section	Disassociation and Deauthentication Attacks	101
Section	Jamming Attacks	102
Section	Radio Frequency Identified (RFID) Attacks	102
Section	Near-Field Communication (NFC) Attacks	102
Section	Initialization Vector (IV) Attacks	103
Figure 4-2	An On-Path Attacks	104
Section	Address Resolution Protocol (ARP) Poisoning Attacks	105
Section	Media Access Control (MAC) Flooding Attacks	106
Section	MAC Cloning Attacks	106
Section	Domain Hijacking Attacks	108
Section	DNS Poisoning Attacks	108
Section	Uniform Resource Locator (URL) Redirection Attacks	110
Section	Domain Reputation	110
Paragraph	Understanding DDoS attacks	111

Key Topic Element	Description	Page Number
Paragraph	Understanding the difference between network and application-targeted DDoS attacks	113
Paragraph	Surveying OT-related attacks	113
Section	Malicious Code or Script Execution Attacks	113

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

evil twin, rogue APs, bluesnarfing, bluejacking, Wi-Fi disassociation attack, jamming, radio frequency identification (RFID), near-field communication (NFC), initialization vector (IV) attack, on-path attack, ARP cache poisoning, MAC flooding attacks, MAC cloning attack, domain hijacking, DNS poisoning, pharming, domain name kiting, URL redirection attacks, domain reputation, Domain Keys Identified Mail (DKIM), Domain-based Message Authentication, Reporting & Conformance (DMARC), distributed denial-of-service (DDoS) attack, botnet, DNS amplification attack, network DDoS attacks, application DDoS attacks, operational technology (OT), PowerShell, Python, Bash, macros, Visual Basic for Applications (VBA)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is a Microsoft scripting language that attackers have used to perform postexploitation activities such as privilege escalation or to enumerate users?
2. What is a standard that was designed to thwart spammers from spoofing your domain to send email?
3. What is a web application vulnerability that could allow an attacker to perform a URL redirection attack?
4. What is the modification of name resolution information that should be in a DNS server's cache?
5. In what type of attack does the attacker sniff the network for valid MAC addresses and then use those MAC addresses to perform other actions?



This chapter covers the following topics related to Objective 1.5 (Explain different threat actors, vectors, and intelligence sources) of the CompTIA Security+ SY0-601 certification exam:

- Actors and threats
  - Advanced persistent threat (APT)
  - Insider threats
  - State actors
  - Hacktivists
  - Script kiddies
  - Criminal syndicates
  - Hackers
    - Authorized
    - Unauthorized
    - Semi-authorized
  - Shadow IT
  - Competitors
- Attributes of actors
  - Internal/external
  - Level of sophistication/capability
  - Resources/funding
  - Intent/motivation
- Vectors
  - Direct access
  - Wireless
  - Email
  - Supply chain
  - Social media
  - Removable media
  - Cloud
- Threat intelligence sources
  - Open source intelligence (OSINT)
  - Closed/proprietary
  - Vulnerability databases
  - Public/private information sharing centers
  - Dark web
  - Indicators of compromise
  - Automated indicator sharing (AIS)
    - Structured threat information eXpression (STIX)/Trusted automated eXchange of indicator information (TAXII)
  - Predictive analysis
  - Threat maps
  - File/code repositories
- Research sources
  - Vendor websites
  - Vulnerability feeds
  - Conferences
  - Academic journals
  - Request for comments (RFC)
  - Local industry groups
  - Social media
  - Threat feeds
  - Adversary tactics, techniques, and procedures (TTP)

Understanding Different Threat Actors, Vectors, and Intelligence Sources

This chapter covers the different types and attributes of threat actors. You also learn what threat intelligence is, what it is used for, and the different types of threat intelligence sources and formats. The chapter concludes with a list of research resources that individuals and organizations can use to understand today's threats and to protect their networks and systems.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 5-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 5-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Actors and Threats	1–2
Attributes of Actors	3–4
Attack Vectors	5–6
Threat Intelligence and Threat Intelligence Sources	7–8
Research Sources	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following threat actors typically aims to cause denial-of-service conditions or deface websites due to a political or social belief?
   a. Hacktivists
   b. Semi-authorized hackers
   c. Script kiddies
   d. Nation state actors

2. Which of the following threat actors is typically motivated by money? (Choose the best answer.)
   a. Authorized hackers
   b. Criminal syndicates and organized crime
   c. Exploit groups
   d. None of these answers are correct.

3. Which of the following attributes make threat actors different?
   a. Level of sophistication
   b. Resources
   c. Funding
   d. All of these answers are correct.

4. Which incident response concept is designed to represent a cybersecurity incident and is made up of four parts?
   a. FIRST
   b. InfraGard
   c. Diamond Model of Intrusion
   d. All of these answers are correct.

5. Which of the following is an effective attack vector where an attacker could modify or tamper hardware or software from a vendor to perform mass compromise attacks?
   a. Supply chain
   b. Removable media
   c. Wireless
   d. Direct access

6. Which of the following elements could an attacker leverage to perform a cloud-based attack?
   a. Misconfigured VMs
   b. Unpatched applications and operating systems
   c. Misconfigured storage buckets
   d. All of these answers are correct.

7. Which term refers to the knowledge about an existing or emerging threat to assets, including networks and systems?
   a. Threat intelligence
   b. Threat feed
   c. Threat model
   d. None of these answers are correct.

8. Which technique is used when you leverage public information from DNS records, social media sites, websites, search engines, and other sources for reconnaissance?
   a. OSINT
   b. Threat maps
   c. Threat models
   d. Threat analysis

9. Which of the following could be threat and vulnerability research sources?
   a. Vendor websites
   b. Threat feeds
   c. Vulnerability feeds
   d. All of these answers are correct.

10. Which framework was created by MITRE to describe adversary tactics and techniques?
    a. InfraGard
    b. CVE
    c. ATT&CK
    d. CWE

Foundation Topics

Actors and Threats

Key Topic

It's important to understand the types of attackers that you might encounter and their characteristics and personality traits. Keep in mind that these "actors" can find ways to access systems and networks just by searching the Internet. Some attacks can be perpetrated by people with little knowledge of computers and technology. However, other actors have increased knowledge, and with knowledge comes power—the type of power you want to be ready for in advance. However, all levels of attackers can be dangerous. Let's look at several of the actors now.

The first of these personas is the ***script kiddie***. This unsophisticated individual has little or no skills when it comes to technology. This person uses code that was written by others and is freely accessible on the Internet. For example, a script kiddie might copy a malicious PHP script directly from one website to another; only the knowledge of "copy and paste" is required. This derogatory term is often associated with juveniles. Though the processes they use are simple, script kiddies can knowingly (and even unwittingly) inflict incredible amounts of damage to insecure systems. These people almost never have internal knowledge of a system and typically have a limited amount of resources, so the amount, sophistication, and extent of their attacks are constrained. They are often thrill-seekers and are motivated by a need to increase their reputation in the script kiddie world.

Next is the ***hacktivist***—a combination of the terms *hack* and *activist*. As with the term *hacker*, the name *hacktivist* is often applied to different kinds of activities—from hacking for social change, to hacking to promote political agendas, to full-blown cyberterrorism. Due to the ambiguity of the term, a hacktivist could be inside a company or attack from the outside and will have varying amounts of resources and funding. However, the hacktivist is usually far more competent than a script kiddie.

Cybercriminals might work on their own, or they might be part of ***criminal syndicates*** and ***organized crime***—a centralized enterprise run by people motivated mainly by money. Individuals who are part of an organized crime group are often well funded and can have a high level of sophistication.

Individuals might also carry out cyber attacks for governments and nation states. In this case, a government—and its processes—is known as an ***advanced persistent threat (APT)***, though this term can also refer to the set of computer-attacking processes themselves. Often, an APT entity has the highest level of resources, including open-source intelligence (OSINT) *and* covert sources of intelligence. This, coupled with extreme motivation, makes the APT one of the most dangerous foes.

In many cases, APTs are tied to nation ***state actors***. It is not a taboo anymore that governments also use cyber techniques as a method of warfare. State actors can use adversarial techniques to steal intellectual property, cause disruption, and potentially compromise critical infrastructure.

NOTE OSINT is covered in more detail later in this chapter.

Companies should be prepared for insiders and competitors as well. In fact, they might be one and the same. An insider threat can be one of the deadliest to servers and networks, especially if the person making this threat is a system administrator or has security clearance at the organization. Not all "insiders" may be malicious, however. Think about the risks that ***shadow IT*** may introduce in a company. In shadow IT an employee or a group of employees use IT systems, network devices, software, applications, and services without the approval of the corporate IT department. For example, an engineer may just deploy a physical or virtual server and host his or her own application in it. The engineer may not patch or monitor that system often for any security vulnerabilities. Consequently, threat actors may easily compromise the vulnerable system or application. This problem has grown exponentially in recent years because more companies and individuals are adopting cloud-based applications and services. It is very easy to "spin off" a new virtual machine (VM), container, or storage element in the cloud. Employees often put sensitive corporate information in cloud services that are often not closely monitored, introducing a significant risk to the corporation.

You should know these actors and their motivations, but in the end, it doesn't matter too much what you call the attackers. Be ready to prevent the basic attacks and the advanced attacks and treat them all with respect.

Not all hackers are malicious, however. That's right! There are different types of hackers. Different organizations use various names, but some of the common labels include the following:

- An ***authorized hacker*** (white hat) is also known an ethical hacker; this term describes those who use their powers for good. They are often granted permission to hack.
- An ***unauthorized hacker*** (black hat) is a malicious hacker who may have a wide range of skill, but this category is used to describe those hackers involved with criminal activities or malicious intent.
- A ***semi-authorized*** (gray hat) ***hacker*** falls somewhere in the middle of authorized and unauthorized. This person doesn't typically have malicious intent but often runs afoul of ethical standards and principles.

Attributes of Threat Actors

Key Topic

Now that you have learned about the different types of threat actors, let's try to understand their attributes, ***intent***, and ***motivation***. The intent and motivation, in most cases, differ depending on the type of attacker. For example, the motivation for hacktivists is typically around protesting against something that they do not agree with (a political or religious belief, human rights, etc.). Hacktivists tend to use disruption techniques, perform denial-of-service (DoS) attacks, or deface websites. The intent of criminals is typically to make money. Nation state actors can also have different motivations and intent. Some nation state actors just want to steal intellectual property from another nation, government, or a corporation; whereas others want to cause disruption and bring down critical infrastructure. Another example is an ***internal actor*** (insider) who in some cases could be motivated by money being offered from an ***external actor*** (either to steal information or to cause disruption). The internal actor could also be a disgruntled employee who may leave a backdoor either to later access confidential information or to trigger a denial-of-service condition on critical systems and applications.

Another difference between the different types of threat actors is the level of sophistication/capability and the resources/funding. Obviously, a disgruntled employee, a hacktivist, or an amateur script kiddie will not have the resources and funding that nation state actors may have at their disposal. Criminal organizations in some cases may have more resources, money, and access to sophisticated exploits than hacktivists.

Active intrusions start with an adversary who is targeting a victim. The adversary will use various capabilities along some form of infrastructure to launch an attack against the victim. Capabilities can be various forms of tools, techniques, and procedures, while the infrastructure is what connects the adversary and victim.

Attack Vectors

Key Topic

A threat actor may leverage different attack vectors:

- **Direct access:** Direct network access or physical access to a device provides one type of attack vector.
- **Wireless:** This attack vector includes hijacking wireless connections, rogue wireless devices, evil twins, and the attacks you learned in Chapter 4, "Analyzing Potential Indicators Associated with Network Attacks."
- **Email:** Email attack vectors include phishing and spear phishing emails with malicious attachments or malicious links.

- **Supply chain:** One of the most effective attacks for mass compromise is to attack the supply chain of a vendor to tamper with hardware and/or software. This tampering might occur in-house or earlier, while in transit through the manufacturing supply chain.
- **Social media:** This attack vector includes leveraging social media platforms for reconnaissance or to launch social engineering attacks.
- **Removable media:** Many attackers have successfully compromised victim systems by just leaving USB sticks (sometimes referred to as USB keys or USB pen drives) unattended or placing them in strategic locations. Often users think that the devices are lost and insert them into their systems to figure out who to return the devices to; before they know it, they may be downloading and installing malware. Plugging in that USB stick you found lying on the street outside your office could lead to a security breach.
- **Cloud:** Attackers can leverage misconfigured and insecure cloud deployments including unpatched applications, operating systems, and storage buckets.

Threat Intelligence and Threat Intelligence Sources

Key Topic

Threat intelligence refers to knowledge about an existing or emerging threat to assets, including networks and systems. Threat intelligence includes context, mechanisms, ***indicators of compromise*** (IoCs), implications, and actionable advice. This type of intelligence includes specifics on the tactics, techniques, and procedures of these adversaries. The primary purpose of threat intelligence is to inform business decisions regarding the risks and implications associated with threats.

Converting these definitions into common language could translate to threat intelligence being evidence-based knowledge of the capabilities of internal and external threat actors. This type of data can be beneficial for the security operations center (SOC) of any organization. Threat intelligence extends cybersecurity awareness beyond the internal network by consuming intelligence from other sources Internet-wide related to possible threats to you or your organization. For instance, you can learn about threats that have impacted different external organizations. Subsequently, you can proactively prepare rather than react after the threat is seen against your network. One service that threat intelligence platforms typically provide is an enrichment data feed.

During the investigation and resolution of a security incident, you may also need to communicate with outside parties regarding the incident. Examples include but are not limited to contacting law enforcement, fielding media inquiries, seeking external expertise, and working with Internet service providers (ISPs), the vendor of your

hardware and software products, threat intelligence vendor feeds, coordination centers, and members of other incident response teams.

You can also share relevant incident IoC information and other threat intelligence with industry peers. There are ***public*** and ***private information sharing centers***. A good example of information-sharing communities includes the Financial Services Information Sharing and Analysis Center (FS-ISAC).

TIP ***Information Sharing and Analysis Centers (ISACs)*** are private-sector critical infrastructure organizations and government institutions that collaborate and share information between each other. ISACs exist for different industry sectors. Examples include automotive, aviation, communications, IT, natural gas, elections, electricity, financial services, health-care, and many other ISACs. You can learn more about ISACs at www.nationalisacs.org.

A cybersecurity incident response plan should account for these types of interactions with outside entities. It should also include information about how to interact with your organization's public relations department, legal department, and upper management. You should also get their buy-in when sharing information with outside parties to minimize the risk of information leakage. In other words, you should avoid leaking sensitive information regarding security incidents with unauthorized parties. These actions could potentially lead to additional disruption and financial loss. You should also maintain a list of all the contacts at those external entities, including a detailed list of all external communications for liability and evidentiary purposes.

You can also subscribe to commercial and open-source threat intelligence feeds. Some of these feeds could be *closed* or *proprietary*. In cybersecurity, there is a concept called ***open-source intelligence (OSINT)***. OSINT applies to offensive security (ethical hacking/penetration testing) and defensive security. In offensive security, OSINT enables you to leverage public information from DNS records, social media sites, websites, search engines, and other sources for reconnaissance—in other words, to obtain information about a targeted individual or an organization. When it comes to threat intelligence, OSINT refers to public and free sources of threat intelligence.

TIP The following GitHub repository includes numerous references about OSINT tools and sources: https://github.com/The-Art-of-Hacking/h4cker/tree/master/osint.

You also could leverage other sources to obtain information about threats, including the ***dark web***. The "deep web" is the part of the Internet that has not been indexed by search engines like Google or Bing. The dark web is just a subset of what is considered the "deep web" and is the place where many threat actors perform

malicious activities such as selling stolen credit card numbers, health records, and other personal information. Criminals also sell drugs, guns, counterfeit equipment, and exploits. In some cases, security professionals go to the dark web to perform research and try to find different threats and exploits that could affect their organization. This task is easier said than done. This is why several companies sell "dark web monitoring" and threat intelligence services.

As a security professional, you also have to pay attention to ***vulnerability databases*** and feeds. One of the most popular vulnerability databases is the United States National Vulnerability Database (NVD), part of the National Institute of Standards and Technology (NIST). You can access the NVD at https://nvd.nist.gov.

The NVD provides information about vulnerabilities disclosed worldwide and is assigned the Common Vulnerability and Exposure (CVE) identifiers. CVE is a standard created by MITRE (www.mitre.org) that provides a mechanism to assign an identifier to vulnerabilities so that you can correlate the reports of those vulnerabilities among sites, tools, and feeds. You can obtain additional information about CVE at https://cve.mitre.org.

When it comes to threat intelligence, several standards and formats allow you to perform ***automated indicator sharing (AIS)*** within your organization or with industry peers. The next section covers the most popular standards and formats for threat intelligence exchange.

Key Topic Structured Threat Information eXpression (STIX) and the Trusted Automated eXchange of Indicator Information (TAXII)

The ***Structured Threat Information eXpression (STIX)*** is an OASIS standard designed for sharing of cyber-attack information. STIX details can contain data such as the IP addresses or domain names of command-and-control servers (often referred to C2 or CnC), malware hashes, and so on. STIX was originally developed by MITRE and is now maintained by OASIS.

The ***Trusted Automated eXchange of Indicator Information (TAXII)*** is also an OASIS standard that provides an open transport mechanism that standardizes the automated exchange of cyber-threat information. TAXII was originally developed by MITRE and is now maintained by OASIS.

NOTE You can access the STIX and TAXII specification details at https://oasis-open.github.io/cti-documentation.

In short, STIX is the actual threat intelligence document in JSON format and TAXII is the transport mechanism.

Example 5-1 provides a STIX threat intelligence document.

Example 5-1 A STIX Threat Intelligence Document

```
{
    "type": "bundle",
    "id": "bundle--aaaa-bbbb-cccc-1234-123123123",
    "objects": [
        {
            "type": "indicator",
            "spec_version": "2.1",
            "id": " indicator-aaabbbcc-b0123-cbc ",
            "created": "2020-10-21T11:09:08.123Z",
            "modified": "2020-10-21T11:09:08.123Z",
            "name": "Windows Trojan Trojan/Win32.PornoAsset.R39927",
            "description": "This file hash indicates that a sample of
Trojan/Win32.PornoAsset.R39927 is present.",
            "indicator_types": [
                "malicious-activity"
            ],
            "pattern": "[file:hashes.'SHA-256' =
'0f1f19244fcc11818083aa1f943bbead338f89e046b8a57a50ec7cf48b62496f']",
            "pattern_type": "stix",
            "valid_from": "2020-03-11T08:00:00Z"
        },
        {
            "type": "malware",
            "spec_version": "2.1",
            "id": " malware--aaaa-bbb-cccc-dddd-eeee ",
            "created": "2020-10-21T11:09:08.123Z",
            "modified": "2020-10-21T11:09:08.123Z",
            "name": "Poison Ivy",
            "malware_types": [
                "remote-access-trojan"
            ],
            "is_family": false
        },
        {
            "type": "relationship",
            "spec_version": "2.1",
            "id": "relationship--12344-4334-1253-6644-b2455",
            "created": "2020-10-21T11:09:08.123Z",
```

```
            "modified": "2020-10-21T11:09:08.123Z",
            "relationship_type": "indicates",
            "source_ref": "indicator-aaabbbcc-b0123-cbc",
            "target_ref": "malware--aaaa-bbb-cccc-dddd-eeee"
        }
    ]
}
```

Key Topic

Predictive analysis and predictive analytics can help discover security threats in your network. Machine-learning-powered solutions have evolved during the last few years, providing benefits that extend beyond signature-based intrusion detection and prevention. Many organizations have to handle large volumes of data from multiple security devices and sources. This makes it almost impossible for individuals to keep up and analyze. As described in Chapter 2, "Analyzing Potential Indicators to Determine the Type of Attack," machine learning (ML) provides a way to quickly analyze security logs and data to perform advanced calculations on it in near real time and help respond to security incidents.

Several organizations have created ***threat maps*** to try to visualize the types of attacks and threats in their network and across the Internet. A cyber-threat map (also known as an attack map) is a real-time map of security attacks happening at any given time. These maps are visualizations from data aggregated from many different sensors around the world. Threat maps help illustrate how prevalent cyber attacks are.

Key Topic

Another source of threat and security information is ***file and code repositories***. Threat actors often release exploits and breached data to file repositories like https://pastebin.com and code repositories like GitLab and GitHub.

Research Sources

Key Topic

Threat and vulnerability research sources include vendor websites, threat feeds, and vulnerability feeds from different providers, including security advisories from different providers. For example, the following link leads to the list of security advisories in Palo Alto products: https://security.paloaltonetworks.com. Likewise, the following link leads to the security advisories from Cisco: https://www.cisco.com/go/psirt. Most mature vendors disclose vulnerabilities affecting their products and services in a transparent way. Several commercial vendors also sell subscriptions to threat and vulnerability feeds.

You can learn and research new threats and vulnerabilities at security conferences, in academic journals, and from local industry groups. Examples of local industry groups include InfraGard (a collaborative effort between the FBI and the private sector), OWASP community chapters, DEF CON groups, and other local groups. Even ***social media*** (like Twitter) can be a source of security threat information.

You can also obtain deep technical information about threats and protocol deficiencies in the Internet Engineering Task Force (IETF) ***request for comments (RFC)*** drafts and publications (www.ietf.org).

Key Topic

The MITRE ATT&CK Framework

Threat hunters assume that an attacker has already compromised the network. Consequently, they need to come up with a hypothesis of what is compromised and how an adversary could have performed the attack. For the threat hunting to be successful, hunters need to be aware of the ***adversary tactics, techniques, and procedures (TTPs)*** that modern attackers use. This is why many organizations use MITRE's ***ATT&CK*** framework to be able to learn about the tactics and techniques of adversaries. You will learn more about how MITRE's ATT&CK can be used in threat hunting.

Threat hunting is not a new concept. Many organizations have performed threat hunting for a long time. However, in the last decade many organizations have recognized that they either have to implement a threat-hunting program or enhance their existing program to better defend their organization.

You can use the MITRE's ATT&CK framework to learn about the tactics and techniques that attackers used during their campaigns. The information in ATT&CK can be extremely useful for threat hunting.

ATT&CK (https://attack.mitre.org) is a collection of different matrices of tactics and techniques. The ATT&CK Matrix for Enterprise includes the tactics and techniques that adversaries use while preparing for an attack, including gathering of information (open-source intelligence, technical and people weakness identification, and more), and every single aspect of exploitation, post-exploitation, command and control, and other actions performed by the attacker. MITRE also provides a list of software (tools and malware) that adversaries use to carry out their attacks. You can access this list at https://attack.mitre.org/software.

You can also obtain detailed information about a specific tool or type of malware used by attackers for different purposes. For example, Figure 5-1 shows the well-known malicious memory-scraping and credential-dumping tool called Mimikatz in the MITRE ATT&CK Navigator.

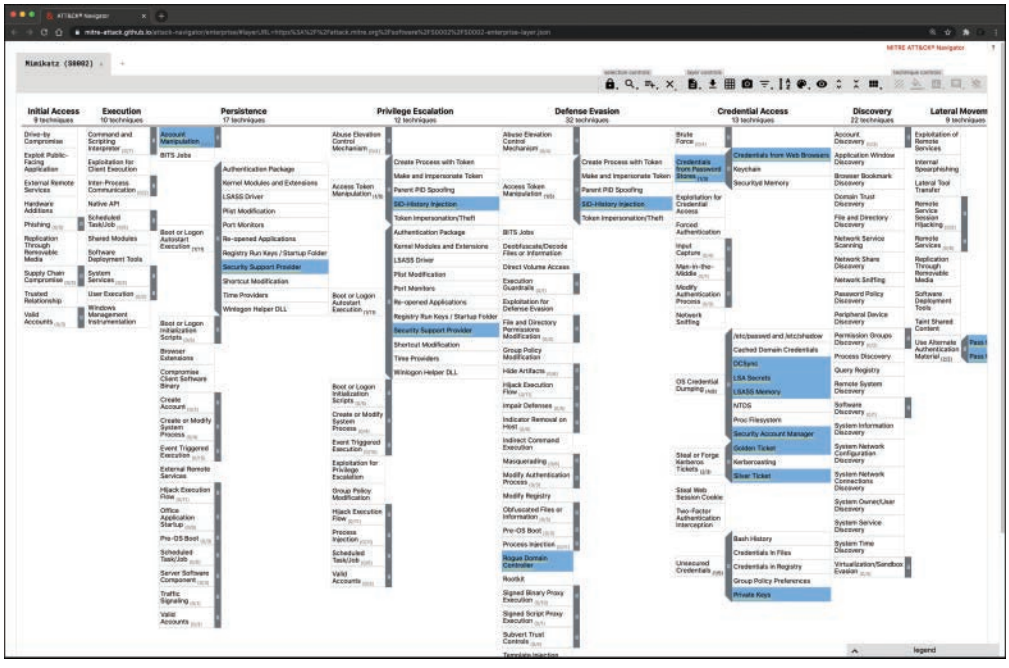


FIGURE 5-1 The MITRE ATT&CK Navigator

As illustrated in Figure 5-1, the MITRE ATT&CK Navigator shows all the different tactics and techniques where Mimikatz has been used in many different attacks by a multitude of attackers.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 5-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 5-2 Key Topics for Chapter 5

Key Topic Element	Description	Page Number
Section	Actors and Threats	120
Paragraph	Understanding the attributes, intent, and motivations of threat actors	122
Section	Attack Vectors	122
Paragraph	Understanding threat intelligence and indicators of compromise	123
Section	Structured Threat Information eXpression (STIX) and the Trusted Automated eXchange of Indicator Information (TAXII)	125
Paragraph	Understanding predictive analysis	127
Paragraph	Identifying threat and security information in file and code repositories	127
Paragraph	Surveying threat and vulnerability research sources such as vendor websites, threat feeds, and vulnerability feeds	127
Section	The MITRE ATT&CK Framework	128

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

script kiddie, hacktivist, criminal syndicates, organized crime, advanced persistent threat (APT), state actors, shadow IT, authorized hacker, unauthorized hacker, semi-authorized hacker, intent, motivation, internal actor, external actor, threat intelligence, indicators of compromise, public information sharing centers, private information sharing centers, Information Sharing and Analysis Centers (ISACs), open-source intelligence (OSINT), dark web, vulnerability databases, automated indicator sharing (AIS), Structured Threat Information eXpression (STIX), Trusted Automated eXchange of Indicator Information (TAXII), predictive analysis, threat maps, file and code repositories, social media, request for comments (RFC), adversary tactics, techniques, and procedures (TTPs), ATT&CK

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What includes the tactics and techniques that adversaries use while preparing for an attack, including gathering of information (open-source intelligence, technical and people weakness identification, and more)?
2. What standard was designed to document threat intelligence in a machine-readable format?
3. What standard was designed as a transport mechanism of threat intelligence and to perform automated indicator sharing (AIS)?
4. What is the vulnerability database maintained by NIST?
5. What is the attack vector where evil twins are used?
6. You were hired to perform a penetration test against three different applications for a large enterprise. You are considered a(n) _____ hacker.
7. You are hired to investigate a cyber attack. Your customer provided different types of information collected from compromised systems such as malware hashes and the IP address of a potential command and control (C2). What are these elements often called?
8. What name is often used to describe a government or state-sponsored persistent and sophisticated attack?
9. What is the term used when an employee or a group of employees use IT systems, network devices, software, applications, and services without the approval of the corporate IT department?



This chapter covers the following topics related to Objective 1.6 (Explain the security concerns associated with various types of vulnerabilities) of the CompTIA Security+ SY0-601 certification exam:

- Cloud-based vs. on-premises vulnerabilities
- Zero-day
- Weak configurations
  - Open permissions
  - Unsecure root accounts
  - Errors
  - Weak encryption
  - Unsecure protocols
  - Default settings
  - Open ports and services
- Third-party risks
  - Vendor management
    - System integration
    - Lack of vendor support
  - Supply chain
  - Outsourced code development
  - Data storage
- Improper or weak patch management
  - Firmware
  - Operating system (OS)
  - Applications
- Legacy platforms
- Impacts
  - Data loss
  - Data breaches
  - Data exfiltration
  - Identity theft
  - Financial
  - Reputation
  - Availability loss

CHAPTER 6

Understanding the Security Concerns Associated with Various Types of Vulnerabilities

This chapter starts with an overview of cloud-based versus on-premises vulnerabilities. You then learn about zero-day vulnerabilities and related attacks. Weak configurations (such as open permissions, unsecure root accounts, errors, weak encryption, unsecure protocols, default settings, and open ports and services) can allow attackers to breach a network and compromise your systems. This chapter covers the details of those common weaknesses. In this chapter, you also learn about third-party risks, the consequences of improper or weak patch management, and the risk of running legacy platforms that are often vulnerable to attacks. The chapter concludes with an overview of the different types of impact categories when a cybersecurity attack or breach occurs.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 6-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 6-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Cloud-based vs. On-premises Vulnerabilities	1–2
Zero-day Vulnerabilities	3
Weak Configurations	4–5
Third-party Risks	6
Improper or Weak Patch Management	7
Legacy Platforms	8
The Impact of Cybersecurity Attacks and Breaches	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which cloud service model offers computer networking, storage, load balancing, routing, and VM hosting?
   a. IaaS
   b. PaaS
   c. SaaS
   d. None of these answers are correct.

2. Which of the following is a security concern in cloud deployments?
   a. Encryption
   b. Authentication methods
   c. Identity management
   d. All of these answers are correct.

3. Which type of vulnerability is disclosed by an individual or exploited by an attacker before the creator of the software can create a patch to fix the underlying issue?
   a. Cross-site scripting
   b. Zero-day
   c. Information disclosure
   d. None of these answers are correct.

4. Which of the following can be considered a weak configuration that can allow attackers to perform an attack and compromise systems?
   a. Default settings and passwords
   b. Weak encryption
   c. Open permissions
   d. All of these answers are correct.

5. Which of the following is a weak protocol that should be avoided?
   a. HTTP without encryption
   b. Telnet
   c. FTP without encryption
   d. All of these answers are correct.

6. Which of the following should be considered when assessing third-party risks?
   a. Vendor management
   b. Supply chain
   c. Outsourced code development
   d. All of these answers are correct.

7. In Windows, what is a broadly released fix for a product-specific security-related vulnerability?
   a. Security update
   b. Service pack
   c. Threat model
   d. None of these answers are correct.

8. Which of the following is a disadvantage of running legacy platforms and products?
   a. They are often affected by security vulnerabilities.
   b. They do not have modern security features.
   c. When a device is past the last day of support, vendors will not investigate or patch security vulnerabilities in those devices.
   d. All of these answers are correct.

9. Which of the following could be categorized as different types of negative impact that a security breach could have in a corporation?
   a. Financial
   b. Reputation
   c. Availability loss
   d. All of these answers are correct.

10. Which of the following can be used by attackers to obfuscate their tactics when exfiltrating data from their target (victim)?
   a. Encryption
   b. Tunneling over a known protocol like DNS
   c. Encoding
   d. All of these answers are correct.

Foundation Topics

Cloud-based vs. On-premises Vulnerabilities

As time moves forward, more and more organizations are transferring some or all of their server and network resources to the cloud. Doing so creates many potential hazards and vulnerabilities that must be addressed by the security administrator and by the cloud provider. Top among these concerns are the servers, where all data is stored and accessed. Servers of all types should be hardened and protected from a variety of attacks in an effort to keep the integrity of data from being compromised. However, data must also be available. Consequently, you, as the security administrator, must strike a balance between security and availability. In this section, we discuss cloud-based threats as well as server vulnerabilities and how to combat them effectively.

Up until now we have focused on the individual computer system. Let's expand our security perimeter to include networks. Network design is extremely important in a secure computing environment. The elements that you include in your design can help to defend against many different types of network attacks. Being able to identify these network threats is the next step in securing your network. If you apply the strategies and defense mechanisms included in this chapter, you should be able to stave off most network-based assaults. Maintaining the security of the servers and network infrastructure of an organization is your job as the security administrator, but with the inclusion of the cloud, the areas of responsibility might vary. These responsibilities depend on how much of the cloud is provided by a third party and how much of the cloud is held privately within the organization's domain. Whether dealing with cloud providers, onsite cloud-based resources, locally owned servers and networks, or a mixture of all of them, you have a lot of duties and must understand not only security but how computer networking really functions. The CompTIA Security+ exam assumes that you have a working knowledge of networks and that you have the CompTIA Network+ certification or commensurate experience. Hereafter, this book will work within that mindset and will refer directly to the security side of things as it progresses. So, put on your networking hat and let's begin with network design.

Historically, the term *cloud* was just a name for the Internet—anything beyond your network that you as a user couldn't see. Technically speaking, the cloud was the area of the telephone company's infrastructure; it was everything between one organization's demarcation point and the demarcation point of another organization. It included central offices, switching offices, telephone poles, circuit-switching devices, packet assemblers/disassemblers (PADs), packet-switching exchanges (PSEs), and so on. In fact, all these things, and much more, are still part of the cloud, in the

technical sense. Back in the day, this term was used only by telecommunications professionals and network engineers.

Today, the term *cloud* has a somewhat different meaning. Almost everyone has heard of it and probably used it to some extent. It is used heavily in marketing, and the meaning is less technical and more service-oriented than it used to be. It takes the place of most intranets and extranets that had existed for decades before its emergence.

In ***on-premises*** environments, physical servers and virtual machines control the sending and receiving of all kinds of data over the network, including websites, email and text messaging, and data stored as single files and in database format. A great many of these "servers" are now virtual in the cloud, with more moving there every day. And the cloud, however an organization connects to it, is all about networking. Therefore, the cloud, virtualization, and servers in general are all thoroughly intertwined.

Key Topic

Cloud computing can be defined as a way of offering on-demand services that extend the capabilities of a person's computer or an organization's network. These might be free services, such as personal browser-based email from various providers, or they could be offered on a pay-per-use basis, such as services that offer data access, data storage, infrastructure, and online gaming. A network connection of some sort is required to make the connection to the cloud and gain access to these services in real time.

Some of the benefits to an organization using cloud-based services include lowered cost, less administration and maintenance, more reliability, increased scalability, and possible increased performance. A basic example of a cloud-based service would be browser-based email. A small business with few employees definitely needs email, but it can't afford the costs of an email server and perhaps does not want to have its own hosted domain and the costs and work that go along with that. By connecting to a free web browser–based service, a small business can obtain near unlimited email, contacts, and calendar solutions. But, there is no administrative control, and there are some security concerns, which we discuss shortly.

Cloud computing services are generally broken down into several categories of services:

Key Topic

- ***Software as a service (SaaS)*:** This is the most commonly used and recognized of the three categories designed to provide a complete packaged solution. The software is rented out to the user. The service is usually provided through some type of front end or web portal. While the end user is free to use the service from anywhere, the company pays a per-use fee. Examples of SaaS offerings are Office 365, Gmail, Webex, Zoom, Dropbox, and Google Drive.

NOTE Often compared to SaaS is the application service provider (ASP) model. SaaS typically offers a generalized service to many users. However, an ASP typically delivers a service (perhaps a single application) to a small number of users.

- ***Infrastructure as a service (IaaS)*:** This service offers computer networking, storage, load balancing, routing, and VM hosting. More and more organizations are seeing the benefits of offloading some of their networking infrastructure to the cloud.
- ***Platform as a service (PaaS)*:** This service provides various software solutions to organizations, especially the ability to develop applications in a virtual environment without the cost or administration of a physical platform. PaaS is used for easy-to-configure operating systems and on-demand computing. Often, this utilizes IaaS as well for an underlying infrastructure to the platform. Cloud-based virtual desktop environments (VDEs) and virtual desktop infrastructures (VDIs) are often considered to be part of this service but also can be part of IaaS.
- **Security as a service (SECaaS):** In this service a large service provider integrates its security services into the company's or customer's existing infrastructure. The concept is that the service provider can provide the security more efficiently and more cost effectively than a company can, especially if it has a limited IT staff or budget. The Cloud Security Alliance (CSA) defines various categories to help businesses implement and understand SECaaS, including encryption, data loss prevention (DLP), continuous monitoring, business continuity and disaster recovery (BCDR), vulnerability scanning, and much more.

NOTE Periodically, new services will arrive, such as monitoring as a service (MaaS)—a framework that facilitates the deployment of monitoring within the cloud in a continuous fashion. There are many types of cloud-based services. If they don't fall into the preceding list, then they often fall under the category "anything as a service" (XaaS).

A cloud service provider (CSP) might offer one or more of these services. The National Institute of Standards and Technology (NIST) provides a great resource explaining the different cloud service models in Special Publication (SP) 800-145 (The NIST Definition of Cloud Computing). However, they are summarized here. There are different types of clouds used by organizations: public, private, hybrid, and community.

- ***Public cloud***: In this type, a service provider offers applications and storage space to the general public over the Internet. Examples include free, web-based email services and pay-as-you-go business-class services. The main benefits of this type of cloud include low (or zero) cost and scalability. Providers of public cloud space include Google Cloud Platform (GCP), Microsoft Azure, and Amazon Web Services (AWS).
- ***Private cloud***: This type is designed with a particular organization in mind. As security administrator, you have more control over the data and infrastructure. A limited number of people have access to the cloud, and they are usually located behind a firewall of some sort in order to gain access to the private cloud. Resources might be provided by a third party or could come from the server room or data center.
- ***Hybrid cloud***: This is a mixture of public and private clouds. Dedicated servers located within the organization and cloud servers from a third party are used together to form the collective network. In these hybrid scenarios, confidential data is usually kept in-house.
- ***Community cloud***: This is another mix of public and private but one where multiple organizations can share the public portion. Community clouds appeal to organizations that usually have a common form of computing and storing of data.

The type of cloud an organization uses will be dictated by its budget, the level of security it requires, and the amount of manpower (or lack thereof) it has to administer its resources. While a private cloud can be very appealing, it is often beyond the ability of an organization, forcing that organization to seek the public or community-based cloud. However, it doesn't matter what type of cloud is used. Resources still have to be secured by someone, and you'll have a hand in that security one way or the other.

Cloud security hinges on the level of control you retain and the types of security controls you implement. When an organization makes a decision to use cloud computing, probably the most important security control concern to administrators is the loss of physical control of the organization's data. A more in-depth list of cloud computing security concerns includes lack of privacy, lack of accountability, improper authentication, lack of administrative control, data sensitivity and integrity problems, data segregation issues, location of data and data recovery problems, malicious insider attacks, bug exploitation, lack of investigative support when there is a problem, and finally, questionable long-term viability. In general, everything that you worry about for your local network and computers! Keep in mind that cloud service providers can be abused as well: attackers often attempt to use providers' infrastructure to launch powerful attacks.

Solutions to these security issues include the following:

- **Complex passwords:** Strong passwords are beyond important; they are critical, as mentioned many times in this text. As of the writing of this book, accepted password schemes include the following:
  - **For general security:** A minimum of 10 characters, including at least one capital letter, one number, and one special character
  - **For confidential data:** A minimum of 15 characters, including a minimum of two each of capital letters, numbers, and special characters

When it comes to the cloud, you might just opt to use the second option for every type of cloud. The reasoning is that public clouds can be insecure (you just don't know), and private clouds will most likely house the most confidential data. To enforce the type of password you want your users to choose, a strong server-based policy is recommended.

However, passwords will not protect your data and systems. Passwords can be stolen and often are reused. This is why multifactor authentication is so important!

- **Powerful authentication methods:** Multifactor authentication can offer a certain amount of defense in depth. In this scenario, if one form of authentication is compromised, the other works as a backup. For example, in addition to a password, a person might be asked for biometric confirmation such as a thumbprint or voice authorization, for an additional PIN, or to swipe a smart card. Multifactor authentication may or may not be physically possible, depending on the cloud environment being used, but if at all possible, it should be considered. An example of a multifactor authentication solution is DUO (https://duo.com). It provides a way to integrate multifactor authentication with many different types of deployments, and users can use a phone app to confirm their identity and authenticate to a user or application.
- **Strong cloud data access policies:** We're talking the who, what, and when. When it comes to public clouds especially, you should specifically define which users have access, exactly which resources they have access to, and when they are allowed to access those resources. Configure strong passwords and consider two-factor authentication. Configure policies from servers that govern the users; for example, use Group Policy objects on a Windows Server domain controller. Audit any and all connected devices and apps. Consider storing different types of data with different services. Some services do better with media files, for example. Remember that cloud storage is not backup. Approach the backing up of data as a separate procedure.
- **Encryption:** Encryption of individual data files, whole disk encryption, digitally signed virtual machine files…the list goes on. Perhaps the most important

is a robust public key infrastructure (PKI), which is discussed further in Chapter 25, "Implementing Public Key Infrastructure." The reason is that many users will access data through a web browser.

- **Standardization of programming:** The way applications are planned, designed, programmed, and run on the cloud should all be standardized from one platform to the next, and from one programmer to the next. Most important is standardized testing in the form of input validation, fuzzing, and known environment, unknown environment, or partially known environment testing.
- **Protection of all the data!:** This includes storage-area networks (SANs), general cloud storage, and the handling of big data (for example, astronomical data). When data is stored in multiple locations, it is easy for some to slip through the cracks. Detailed documentation of what is stored where (and how it is secured) should be kept and updated periodically. As a top-notch security administrator, you don't want your data to be tampered with. Therefore, implementing some cloud-based security controls can be very helpful. For example, consider the following: deterrent controls (prevent the tampering of data), preventive controls (increase the security strength of a system that houses data), corrective controls (reduce the effects of data tampering that has occurred), and detective controls (detect attacks in real time, and have a defense plan that can be immediately carried out).

What else are you, as administrator, trying to protect here? You're concerned with protecting the identity and privacy of your users (especially executives because they are high-profile targets). You need to secure the privacy of credit card numbers and other super-confidential information. You want to secure physical servers that are part of the server room or data center because they might be part of your private cloud. You desire protection of your applications with testing and acceptance procedures. (Keep in mind that these things all need to be done within contractual obligations with any third-party cloud providers.) And finally, you're interested in promoting the availability of your data. After all of your security controls and methods have been implemented, you might find that you have locked out more people than first intended. So, your design plan should contain details that will allow for available data, but in a secure manner.

Customers considering using cloud computing services should ask for transparency—or detailed information about the provider's security. The provider must be in compliance with the organization's security policies; otherwise, the data and software in the cloud become far less secure than the data and software within the customer's own network. This concept, and most of the concepts in the first half of this chapter, should be considered when planning whether to have data, systems, and infrastructure

contained on-premises, in a hosted environment, on the cloud, or in a mix of those. If there is a mix of on-premises infrastructure and cloud-provider infrastructure, a company might consider a ***cloud access security broker (CASB)***—a software tool or service that acts as the gatekeeper between the two, allowing the company to extend the reach of its security policies beyond its internal infrastructure.

Other "Cloud"-based Concerns

Other technologies to watch out for are loosely connected with what can be called cloud technologies. One example is social media. Social media environments can include websites as well as special applications that are loaded directly on to the computer (mobile or desktop), among other ways to connect, both legitimate and illegitimate. People share the darndest things on social media websites, which can easily compromise the security of employees and data. The point? There are several ways to access social media platforms, and it can be difficult for a security administrator to find every website, application, service, and port that is used by social media. In cases such as these, you might consider implementing more allow lists and block/deny lists to safeguard applications so that users are better locked down.

Another thing to watch for is P2P networks. File sharing, gaming, media streaming, and all the world is apparently available to a user—if the user knows where to look. However, P2P often comes with a price: malware and potential system infiltration. By the latter, I mean that computers can become unwilling participants in the sharing of data on a P2P network. This is one example in which the cloud invades client computers, often without the user's consent. Access to file sharing, P2P, and torrents also needs a permanent "padlock."

Then there's the dark web. Where Batman stores his data… No, the dark web is another type of P2P (often referred to as an F2F, meaning friends to friends) that creates connections between trusted peers (unlike most other P2Ps) but uses non-standard ports and protocols. This makes it a bit more difficult to detect. The dark web is certainly the safe haven of illegal activities because they are designed specifically to resist surveillance. Computers that are part of an administrator's network, and more often, virtual machines in the admin's cloud, can be part of the dark web and can easily go undetected by the admin. In some cases, an employee of the organization (or an employee of the cloud provider) might have configured some cloud-based resources to join the dark web. This can have devastating legal consequences if illegal activities are traced to your organization. Thorough checks of cloud-based resources can help to prevent this situation. Also, screening of employees, careful inspection of service-level agreements with cloud providers, and the use of third-party IT auditors can avoid the possibility of dark web connectivity, P2P links, and improper use of social media.

Server Defense

Now we come down to it. Servers are the cornerstone of data. They store it, transfer it, archive it, and allow or disallow access to it. They need super-fast network connections that are monitored and baselined regularly. They require you to configure policies, check logs, and perform audits frequently. They exist in networks both large and small, within public and private clouds, and are often present in virtual fashion. What it all comes down to is that servers contain the data and the services that everyone relies on. So, they are effectively the most important things to secure on your network.

Let's break down five types of servers that are of great importance (in no particular order), and look at some of the threats and vulnerabilities to those servers, and ways to protect them.

File Servers

File server computers store, transfer, migrate, synchronize, and archive files. Really any computer can act as a file server of sorts, but examples of actual server software include Microsoft Windows Server and the various types of Linux server versions (for example, Ubuntu Server or Red Hat Server), not to mention UNIX. File servers are vulnerable to the same types of attacks and malware that typical desktop computers are. To secure file servers (and the rest of the servers on this list), you should employ hardening, updating, antimalware applications, software-based firewalls, hardware-based intrusion detection systems (HIDSs), and encryption, and be sure to monitor the server regularly.

Network Controllers

A network controller is a server that acts as a central repository of user accounts and computer accounts on the network. All users log in to this server. An example would be a Windows Server system that has been promoted to a domain controller (running Active Directory). In addition to the attacks mentioned for file servers, a domain controller can be the victim of LDAP injection. It also has Kerberos vulnerabilities, which can ultimately result in privilege escalation or spoofing. LDAP injection can be prevented with proper input validation. But in the specific case of a Windows domain controller, really the only way to keep it protected (aside from the preventive measures mentioned for file servers) is to install specific security software updates for the OS.

Email Servers

Email servers are part of the message server family. Here, a message server is any server that deals with email, faxing, texting, chatting, and so on. For this section let's concentrate strictly on the email server. The most common of these is Microsoft Exchange. An Exchange Server might run POP3, SMTP, and IMAP, and allow for Outlook web-based connections. That's a lot of protocols and ports running. So, it's not surprising to hear some Exchange admins confess that running an email server can be difficult at times, particularly because it is vulnerable to XSS attacks, overflows, DoS/DDoS attacks, SMTP memory exploits, directory traversal attacks, and of course, spam. Bottom line: it has to be patched…a lot.

As an administrator, you need to keep on top of the latest vulnerabilities and attacks and possibly be prepared to shut down or quarantine an email server at a moment's notice. New attacks and exploits are constantly surfacing because email servers are a common and big target with a large attack surface. For spam, a hardware-based spam filter is most effective (such as one from Barracuda), but software-based filters can also help. To protect the integrity and confidentiality of email-based data, you should consider DLP and encryption. Security could also come in the form of secure POP3 and secure SMTP; for more about the specific secure email protocols, see Chapter 7, "Summarizing the Techniques Used in Security Assessments." Also, security can come as encrypted SSL/TLS connections, security posture assessment (SPA), secure VPNs, and other encryption types, especially for web-based connections. Web-based connections can be particularly insecure; great care must be taken to secure these connections. For example, push notification services for mobile devices are quite common. While TLS is normally used as a secure channel for the email connection, text and metadata can at times be sent as clear text. A solution to this is for the operating system to use a symmetrical key to encrypt the data payload. Vendors may or may not do this, so it is up to the email administrator to incorporate this added layer of security, or at least verify that push email providers are implementing it.

Thinking a little outside of the box, you could consider moving away from Microsoft (which is the victim of the most attacks) and toward a Linux solution such as the Java-based SMTP server built into Apache, or with a third-party tool such as Zimbra (or one of many others). These solutions are not foolproof, and still need to be updated, but it is a well-known fact that historically Linux has not been attacked as often as Microsoft (in general), though the difference between the two in the number of attacks experienced has shrunk considerably since the turn of the millennium.

Web Servers

The web server could be the most commonly attacked server of them all. Examples of web servers include Microsoft's Internet Information Services (IIS), Apache HTTP Server (Linux), lighttpd (FreeBSD), Oracle iPlanet Web Server (Oracle), and iPlanet's predecessor Sun Java System Web Server (Sun Microsystems). Web servers in general can be the victim of DoS attacks, overflows, XSS and XSRF, remote code execution (RCE), and various attacks that make use of backdoors. For example, in IIS, if basic authentication is enabled, a backdoor could be created, and attackers could ultimately bypass access restrictions. An IIS administrator must keep up to date with the latest vulnerabilities by reading Microsoft Security Bulletins, such as this one that addresses possible information disclosure: https://technet.microsoft.com/library/security/ms12-073.

In general, as security administrator, you should keep up to date with ***Common Vulnerabilities and Exposures (CVE)*** as maintained by MITRE (https://cve.mitre.org/). The latest CVE listings for applications and operating systems can be found there and at several other websites.

Aside from the usual programmatic solutions to vulnerabilities such as XSS and standard updating and hot patching, you might consider adding and configuring a hardware-based firewall from Cisco, Fortinet, Palo Alto, or other similar company. Of course, HTTPS (be it SSL or, better yet, TLS) can be beneficial if the scenario calls for it. Once a server is secured, you can prove the relative security of the system to users by using an automated vulnerability scanning program (such as Netcraft) that leaves a little image on the web pages stating whether or not the site is secure and when it was scanned or audited.

Apache can be the casualty of many attacks as well, including privilege escalation, code injection, and exploits to the proxy portion of the software. PHP forms and the PHP engine could act as gateways to the Apache web server. Patches to known CVEs should be applied as soon as possible.

NOTE You can find a list of CVEs to Apache HTTP Server (and the corresponding updates) at https://httpd.apache.org/security.

When it comes to Apache web servers, you have to watch out for the web server attack called Darkleech. It takes the form of a malicious Apache module (specifically an injected HTML iframe tag within a PHP file). If loaded on a compromised Apache web server, it can initiate all kinds of attacks and deliver various payloads of malware and ransomware. Though Darkleech is not limited to Apache, the bulk of Darkleech-infected sites have been Apache-based.

NOTE So much for Microsoft being less targeted than Linux. As time moves forward, it seems that no platform is safe. A word to the wise: don't rely on any particular technology because of a reputation, and be sure to update and patch every technology you use.

As far as combatting Darkleech, a webmaster can attempt to query the system for PHP files stored in folders with suspiciously long hexadecimal names. If convenient for the organization, all iframes can be filtered out. Of course, the Apache server should be updated as soon as possible, and if necessary, taken offline while it is repaired. In many cases, this type of web server attack is very hard to detect, and sometimes the only recourse is to rebuild the server (or virtual server image) that hosts the Apache server.

NOTE Another tool that some attackers use is archive.org. This website takes snapshots of many websites over time and stores them. They are accessible to anyone and can give attackers an idea of older (and possibly less secure) pages and scripts that used to run on a web server. It could be that these files and scripts are still located on the web server even though they are no longer used. This is a vulnerability that you should be aware of. Strongly consider removing older unused files and scripts from web servers.

FTP Server

An FTP server can be used to provide basic file access publicly or privately. Examples of FTP servers include the FTP server built into IIS, the Apache FtpServer, and other third-party offerings such as FileZilla Server and Pure-FTPd.

The standard, default FTP server is pretty insecure. It uses well-known ports (20 and 21), doesn't use encryption by default, and has basic username/password authentication. As a result, FTP servers are often the victims of many types of attacks. Examples include bounce attacks (when a person attempts to hijack the FTP service to scan other computers), buffer overflow attempts (when an attacker tries to send an extremely long username or password or filename to trigger the overflow), and attacks on the anonymous account (if utilized).

If the files to be stored on the FTP server are at all confidential, you should consider additional security. You can do so by incorporating FTP software that utilizes secure file transfer protocols such as FTPS or SFTP. Additional security can be provided by using FTP software that uses dynamic assignment of data ports, instead of using port 20 every time. Encryption can prevent most attackers from reading the data

files, even if they are able to get access to them. Of course, if not a public FTP, the anonymous account should be disabled.

There also are other, more sinister attacks lurking, ones that work in conjunction with the web server, which is often on the same computer, or part of the same software suite—for instance, the web shell. There are plenty of variants of the web shell, but here we describe its basic function. The web shell is a program that is installed on a web server by an attacker and is used to remotely access and reconfigure the server without the owner's consent.

Web shells are remote access Trojans (RATs) but are also referred to as backdoors because they offer an alternative way of accessing the website for the attacker. The reason the web shell attack is described here in the FTP section is that it is usually the FTP server that contains the real vulnerability—weak passwords. Once an attacker figures out an administrator password of the FTP server (often through brute-force attempts), the attacker can easily install the web shell and effectively do anything desired to that web server (and/or the FTP server). It seems like a house of cards, and in a way, it is.

How can you prevent this from happening? First, you can increase the password security and change the passwords of all administrator accounts. Second, you can eliminate any unnecessary accounts, especially any superfluous admin accounts and the dreaded anonymous account. Next, you should strongly consider separating the FTP server and web server to two different computers or virtual machines. Finally, you should set up automated scans for web shell scripts (usually PHP files) or have the web server provider do so. If the provider doesn't offer that kind of scanning, you can use a different provider. If a web shell attack is accomplished successfully on a server, you must at the very least search for and delete the original RAT files and at worst reimage the system and restore from backup. This latter option is often necessary if the attacker has had some time to compromise the server. Some organizations have policies that state servers must be reimaged if they are compromised in any way, shape, or form. This solution is a way of starting anew with a clean slate, but it means a lot of configuring. But again, the overall concern here is the complexity of, and the frequency of changing, the password.

That's the short list of servers. There are plenty of others you need to be cognizant of, including DNS servers, application servers, virtualization servers, firewall/proxy servers, database servers, print servers, remote connectivity servers such as RRAS and VPN, and computer telephony integration (CTI) servers. If you are in charge of securing a server, be sure to examine the CVEs and bulletins for that software and be ready to hot-patch the system at a moment's notice. This means having an RDP, VNC, or other remote connection to that specific server ready to go on your desktop so that you can access it quickly.

Zero-day Vulnerabilities

A ***zero-day vulnerability*** is a type of vulnerability that is disclosed by an individual or exploited by an attacker before the creator of the software can create a patch to fix the underlying issue. Attacks leveraging zero-day vulnerabilities can cause damage even after the creator knows of the vulnerability because it may take time to release a patch to prevent the attacks and fix damage caused by them.

Zero-day attacks can be prevented by using newer operating systems that have protection mechanisms and by updating those operating systems. They can also be prevented by using multiple layers of firewalls and by using application approved lists, which allow only known good applications to run. Collectively, these preventive methods are referred to as zero-day protection.

Table 6-2 summarizes programming vulnerabilities/attacks covered so far.

Key Topic

Table 6-2 Programming Vulnerabilities and Attacks

Vulnerability	Description
Backdoor	An attack placed by programmers, knowingly or inadvertently, to bypass normal authentication, and other security mechanisms in place.
Buffer overflow	A vulnerability that occurs when a process stores data outside the memory that the developer intended.
Remote code execution (RCE)	A situation that occurs when an attacker obtains control of a target computer through some sort of vulnerability, gaining the power to execute commands on that remote computer.
Cross-site scripting (XSS)	A vulnerability that exploits the trust a user's browser has in a website through code injection, often in web forms.
Cross-site request forgery (XSRF)	A vulnerability that exploits the trust that a website has in a user's browser, which becomes compromised and transmits unauthorized commands to the website.
Code injection	An attack that occurs when user input in database web forms is not filtered correctly and is executed improperly. SQL injection is a very common example.
Directory traversal	A method of accessing unauthorized parent (or worse, root) directories.
Zero-day	A group of attacks executed on vulnerabilities in software before those vulnerabilities are known to the creator.

The CompTIA Security+ exam objectives don't expect you to be a programmer, but they do expect you to have a basic knowledge of programming languages and methodologies so that you can help to secure applications effectively. I recommend a basic knowledge of programming languages used to build applications, such as Visual Basic, C++, C#, PowerShell, Java, and Python, as well as web-based programming languages such as HTML, JavaScript, and PHP, plus knowledge of database programming languages such as SQL. This foundation knowledge will help you not only on the exam but also when you, as security administrator, have to act as a liaison to the programming team or if you are actually involved in testing an application.

Weak Configurations

Weak configurations can be leveraged by attackers to compromise systems and networks. Often on-premises and cloud-based systems and applications could be deployed without security in mind. Sometimes, these weak configurations are due to lack of security awareness and understanding of common threats.

The following are some of the most prevalent types of weak configurations.

Key Topic

- **Open permissions:** User accounts can be added to individual computers or to networks. For example, a Windows client, Linux computer, or Mac can have multiple users. Larger networks that have a controlling server—for example, a Windows domain controller—enable user accounts that can access one or more computers on the domain. In some cases, users and/or applications are given access (permissions) to resources that they do not need to access (including files, folders, and systems). Applications should be coded and run in such a way as to maintain the principle of least privilege. Users should have access only to what they need. Processes should run with only the bare minimum access needed to complete their functions. However, this can be coupled with separation of privilege, where access to objects depends on more than one condition (for example, authentication plus an encrypted key).
- **Unsecure root accounts:** Out-of-the-box offerings should be as secure as possible. If possible, user password complexity and password aging default policies should be configured by the programmer, not the user. Administrative accounts (root accounts in Linux or administrator in Windows) should be protected at all times.
- **Errors:** Sometimes weak configurations are due to human errors that could lead to security nightmares. The potential for human error always exists.
- **Weak encryption:** Weak encryption or no encryption can be the worst thing that can happen to a wireless or wired network. This situation can occur for several reasons; for example, someone wants to connect an older device or a device that hasn't been updated, and that device can run only a weaker, older

type of encryption. Weak cryptographic implementations can lead to sensitive data being exposed to attackers and could also have a direct impact to privacy. You will learn the basics of cryptographic concepts and encryption in Chapter 16, "Summarizing the Basics of Cryptographic Concepts."

- **Insecure protocols:** Insecure protocols can also lead to the unauthorized exposure of sensitive data and can allow attackers to compromise systems and applications. Protocols such as Telnet, FTP (without encryption), Finger, SSL, old versions of SNMP, and others should be avoided at all times. Table 6-3 shows protocols and related ports, as well as the "more secure options" of some of those protocols.
- **Default settings:** A common adage in the security industry is, "Why do you need hackers if you have default settings and passwords?" Many organizations and individuals leave infrastructure devices such as routers, switches, wireless access points, and even firewalls configured with default passwords. Attackers can easily identify and access systems that use shared default passwords. It is extremely important to always change default manufacturer passwords and restrict network access to critical systems. A lot of manufacturers now require users to change the default passwords during initial setup, but some don't. Attackers can easily obtain default passwords and identify Internet-connected target systems. Passwords can be found in product documentation and compiled lists available on the Internet. Different examples are available in the GitHub repository at https://github.com/The-Art-of-Hacking/h4cker, but dozens of other sites contain default passwords and configurations on the Internet. It is easy to identify devices that have default passwords and that are exposed to the Internet by using search engines such as Shodan (https://www.shodan.io).
- **Open ports and services:** Any unnecessary ports should be closed, and any open ports should be protected and monitored carefully. Although there are 1,024 well-known ports, for the exam you need to know only a handful of them, plus some that are beyond 1,024, as shown in Table 6-3. Remember that these inbound port numbers relate to the applications, services, and protocols that run on a computer, often a server. When it comes to the OSI reference model, the bulk of these protocols are application layer protocols. Examples of these protocols include HTTP, FTP, SMTP, SSH, DHCP, and POP3, and there are many more. Because these are known as application layer protocols, their associated ports are known as *application service ports*. The bulk of Table 6-3 is composed of application service ports. Some of the protocols listed make use of TCP transport layer connections only (for example, HTTP, port 80; and HTTPS, port 443). Some make use of UDP only (for example, SNMP, port 161). Many can use TCP *or* UDP transport mechanisms. Study Table 6-3 carefully now, bookmark it, and refer to it often.

Key Topic

Table 6-3 Ports and Their Associated Protocols

Port Number	Associated Protocol (or Keyword)	TCP/UDP Usage	Secure Version and Port	Usage
21	FTP	TCP	FTPS, port 989/990	Transfers files from host to host.
22	SSH	TCP or UDP		Secure Shell: Remotely administers network devices and systems. Also is used by Secure Copy (SCP) and Secure FTP (SFTP).
23	Telnet	TCP or UDP		Remotely administers network devices (deprecated).
25	SMTP	TCP	SMTP with SSL/TLS, port 465 or 587	Sends email.
49	TACACS+	TCP		Provides remote authentication. Can also use UDP, but TCP is the default. Compare with RADIUS.
53	DNS	TCP or UDP	DNSSEC	Resolves hostnames to IP addresses and vice versa.
69	TFTP	UDP		Is a basic version of FTP.
80	HTTP	TCP	HTTPS (uses SSL/TLS), port 443	Transmits web page data.
88	Kerberos	TCP or UDP		Provides network authentication; uses tickets.
110	POP3	TCP	POP3 with SSL/TLS, port 995	Receives email.
119	NNTP	TCP		Transports Usenet articles.
135	RPC/epmap/dcom-scm	TCP or UDP		Enables you to locate DCOM ports. Also known as RPC (Remote Procedure Call).
137–139	NetBIOS	TCP or UDP		Enables name querying, data sending, and NetBIOS connections.
143	IMAP	TCP	IMAP4 with SSL/TLS, port 993	Retrieves email, with advantages over POP3.
161	SNMP	UDP		Remotely monitors network devices.

Port Number	Associated Protocol (or Keyword)	TCP/UDP Usage	Secure Version and Port	Usage
162	SNMPTRAP	TCP or UDP		Sends Traps and InformRequests to the SNMP Manager on this port.
389	LDAP	TCP or UDP	LDAP over SSL/TLS, port 636	Maintains directories of users and other objects.
445	SMB	TCP		Provides shared access to files and other resources.
514	Syslog	UDP	A secure version (Syslog over TLS) uses TCP as the transport mechanism and port 6514.	Is used for computer message logging, especially for router and firewall logs.
860	iSCSI	TCP		Is an IP-based protocol used for linking data storage facilities. Also uses port 3260 for the iSCSI target.
1433	Ms-sql-s	TCP		Opens queries to Microsoft SQL server.
1701	L2TP	UDP		Is a VPN protocol with no inherent security. Often used with IPsec.
1723	PPTP	TCP or UDP		Is a VPN protocol with built-in security.
1812/1813	RADIUS	UDP		Is an AAA protocol used for authentication (port 1812), authorization, and accounting (port 1813) of users. Also, ports 1645 and 1646.
3225	FCIP	TCP or UDP		Encapsulates Fibre Channel frames within TCP/IP packets. Contrast with Fibre Channel over Ethernet (FCoE), which relies on the data link layer and doesn't rely on TCP/IP directly.
3389	RDP	TCP or UDP		Remotely views and controls other Windows systems.
3868	Diameter	TCP (or SCTP)		Is a protocol that can be used for authentication, authorization, and accounting (AAA). Diameter is an alternative to the RADIUS protocol.

NOTE You can find a complete list of ports and their corresponding protocols at http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml. Not all protocols have set port numbers. For example, the Real-Time Transport Protocol (RTP) and Secure RTP (SRTP) use a pair of port numbers determined by the application that is streaming the audio and video information via RTP. They are selected from a broad range of ports (between 16384 and 32767).

Unnecessary applications and services use valuable hard drive space and processing power. More importantly, they can be vulnerabilities to an operating system. That's why many organizations implement the concept of *least functionality*. This means that an organization configures computers and other information systems to provide only the essential functions. Using this method, you restrict applications, services, ports, and protocols. This control—called CM-7—is described in more detail by NIST at https://nvd.nist.gov/800-53/Rev4/control/CM-7.

The United States Department of Defense describes this concept in DoD instruction 8551.01 at https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/855101p.pdf.

It's this mindset that can help protect systems from threats that are aimed at insecure applications and services. For example, instant messaging programs can be dangerous. They might be fun for the user but usually are not productive in the workplace (to put it nicely); and from a security viewpoint, they often have backdoors that are easily accessible to attackers. Unless they are required by tech support, they should be discouraged and/or disallowed by rules and policies. You should be proactive when it comes to these types of programs. If a user can't install an IM program to a computer, then you will never have to remove it from that system. However, if you do have to remove an application like this, be sure to remove all traces that it ever existed. That is just one example of many, but it can be applied to most superfluous programs.

Other programs you should watch out for are remote control programs. Applications that enable remote control of a computer should be avoided if possible. For example, Remote Desktop Connection is a commonly used Windows-based remote control program. By default, this program uses inbound port 3389, which is well known to attackers—an obvious security threat. Consider using a different port if the program is necessary, and if not, make sure that the program's associated service is turned off and disabled. Check whether any related services need to be disabled as well. Then verify that their inbound ports are no longer functional and that they are closed and secured. Confirm that any shares created by an application are disabled as well. Basically, remove all instances of the application or, if necessary, reimage the computer!

Uninstalling applications on a few computers is feasible, but what if you have a larger network? Say, one with 1,000 computers? You can't expect yourself or your computer techs to go to each and every computer locally and remove applications. That's when centrally administered management systems come into play. Examples include Microsoft's Endpoint Configuration Manager and the variety of mobile device management (MDM) suites available. These programs allow you to manage lots of computers' software, configurations, and policies, all from a local workstation.

Third-party Risks

When an organization is dealing with computer security, a *risk* is the possibility of a malicious attack or other threat causing damage or downtime to a computer system. Generally, this is done by exploiting vulnerabilities in a computer system or network. The more vulnerability, the more risk. Smart organizations are extremely interested in managing vulnerabilities and thereby managing risk. *Risk management* can be defined as the identification, assessment, and prioritization of risks, and the mitigating and monitoring of those risks. Specifically, in relation to computer hardware and software, risk management is also known as *information assurance (IA)*. Many well-known breaches have been caused by attackers leveraging misconfigured or vulnerable systems and applications in third-party vendors (such as a business partner or vendor).

Key Topic

Vendor management is crucial for any organization. Whether you are part of a small company or a global enterprise, third-party vendors have become essential to business operations. These third-party vendors include outsourcing companies, the software and hardware you purchase for your infrastructure, cloud service providers, and more. Risk management should include the evaluation of third-party vendors, along with their security capabilities, and even include security requirements in contract negotiations. In addition, if you are employing contractors from an external vendor, you should always audit the level of access that those contractors have of your resources and data.

Many organizations hire external ***system integration*** companies to deploy new technologies and applications. You must understand your system's integrator security capabilities and understand what risks misconfigurations or vulnerable designs can bring to your organization.

Key Topic

Your organization may also leverage ***outsourced code development***. No software is immune to security vulnerabilities. You must have a way for the vendor to validate and test for any software security vulnerabilities that could be introduced in the outsourced code. You should ask vendors to demonstrate how they test for vulnerabilities including but not limited to static and dynamic analysis, as well as software composition analysis.

Key Topic

Lack of vendor support in current and legacy software and systems also introduces a big risk to your organization. If software and hardware vulnerabilities are not available due to lack of vendor support, those vulnerabilities can be leveraged by attackers to compromise your systems. The risk introduced in legacy platforms and software is covered later in this chapter. You should consistently reassess your vendor processes and support.

In Chapter 5, "Understanding Different Threat Actors, Vectors, and Intelligence Sources," you learned about ***supply chain*** attacks. Supply chain attacks are also called value-chain attacks. One of the most effective attacks for mass compromise is to attack the supply chain of a vendor to tamper with hardware and/or software. This tampering might occur in-house or, previously, while in transit through the manufacturing supply chain. In addition to traditional vendors of software and hardware, cloud offerings are also susceptible to supply chain attacks. Cloud services are an integral part of most businesses nowadays. The compromise of the underlying cloud infrastructure or services can be catastrophic. For example, an attacker implanting a backdoor in all the base images used for your virtual machines in the cloud could lead to massive compromise of your applications and many of their customers.

Key Topic

Insecure ***data storage*** is a major concern of most companies nowadays—not only for data stored in your data center but also in the cloud. Without a doubt, the data needs to be encrypted, but more specifically three types of data: data in use, data at rest, and data in transit. *Data in processing* (also known as data in use) can be described as actively used data undergoing constant change; for example, it could be stored in databases or spreadsheets. *Data at rest* is inactive data that is archived—backed up to tape or otherwise. *Data in transit* (also known as data in motion) is data that crosses the network or data that currently resides in computer memory.

Redundant systems and data backups should also be taken seriously. The whole concept revolves around single points of failure. A *single point of failure* is an element, object, or part of a system that, if it fails, causes the whole system to fail. By implementing redundancy, you can bypass just about any single point of failure. You will learn more about data and system redundancy in Chapter 13, "Implementing Cybersecurity Resilience."

The computer file system used dictates a certain level of security. On Microsoft computers, the best option is to use NTFS, which is more secure, enables logging (oh so important), supports encryption, and has support for a much larger maximum partition size and larger file sizes. Just about the only place where FAT32 and NTFS are on a level playing field is that they support the same number of file formats. By far, then, NTFS is the best option. If a volume uses FAT or FAT32, it can be *converted* to NTFS using the following command:

```
convert volume /FS:NTFS
```

For example, if you want to convert a USB flash drive named M: to NTFS, the syntax would be

```
convert M: /FS:NTFS
```

There are additional options for the **convert** command. To see them, simply type **convert /?** at the Command Prompt. NTFS enables file-level security and tracks permissions within access control lists (ACLs), which are a necessity in today's environment. Most systems today already use NTFS, but you never know about flash-based and other removable media. Using the **chkdsk** command at the Command Prompt or right-clicking the drive in the GUI and selecting Properties can tell you what type of file system it runs.

Generally, the best file system for Linux systems is ext4. It allows for the best and most configurable security. To find out the file system used by your version of Linux, use the **fdisk -l** command or **df -T** command.

System files and folders, by default, are hidden from view to protect a Windows system, but you never know. To permanently configure the system to not show hidden files and folders, navigate to the File Explorer (or Folder) Options dialog box. Then select the View tab, and under Hidden Files and Folders, select the Don't Show Hidden Files, Folders, or Drives radio button. To configure the system to hide protected system files, select the Hide Protected Operating System Files checkbox, located below the radio button previously mentioned. This way, you turn off the ability to view such files and folders as bootmgr and pagefile.sys. You might also need to secure a system by turning off file sharing, which in most versions of Windows can be done within the Network and Sharing Center.

In the past, I have made a bold statement: "Hard disks *will* fail." But it's all too true. It's not a matter of *if*; it's a matter of *when*. By maintaining and hardening the hard disk with various hard disk utilities, you attempt to stave off that dark day as long as possible. You can implement several strategies when maintaining and hardening a hard disk:

Key Topic

- **Remove temporary files:** Temporary files and older files can clog up a hard disk, cause a decrease in performance, and pose a security threat. It is recommended that Disk Cleanup or a similar program be used. Policies can be configured (or written) to run Disk Cleanup every day or at logoff for all the computers on the network.
- **Periodically check system files:** Every once in a while, it's a good idea to verify the integrity of operating system files. You can do a file integrity check in the following ways:
  - With the **chkdsk** command in Windows. This command examines the disk and provides a report. It can also fix some errors with the **/F** option.

- With the **SFC** (System File Checker) command in Windows. This utility checks and, if necessary, replaces protected system files. It can be used to fix problems in the operating system and in other applications such as Internet Explorer. A typical command you might type is **SFC /scannow**. You can use this command if **chkdsk** is not successful at making repairs.
- With the **fsck** command in Linux. This command is used to check and repair a Linux file system. The synopsis of the syntax is **fsck** [**-sAVRTNP**] [**-C** [**fd**]] [**-t fstype**] [**filesys ...**] [**--**] [**fs-specific-options**]. You can find more information about this command at the corresponding man page for **fsck**. A derivative, **e2fsck**, is used to check a Linux ext2fs (second extended file system). Also, you can download open-source data integrity tools for Linux, such as Tripwire.

- **Defragment drives:** Applications and files on hard drives become fragmented over time. For a server, this could be a disaster because the server cannot serve requests in a timely fashion if the drive is too thoroughly fragmented. To defragment the drive, you can use Microsoft's Disk Defragmenter, the command-line **defrag** command, or other third-party programs.
- **Back up data:** Backing up data is critical for a company. It is not enough to rely on a fault-tolerant array. Individual files or the entire system can be backed up to another set of hard drives, to optical discs, to tape, or to the cloud. Microsoft domain controllers' Active Directory databases are particularly susceptible to attack; the system state for these operating systems should be backed up in case the server fails and the Active Directory needs to be recovered in the future.
- **Use restoration techniques:** In Windows, restore points should be created on a regular basis for servers and workstations. The System Restore utility (rstrui.exe) can fix issues caused by defective hardware or software by reverting back to an earlier time. Registry changes made by hardware or software are reversed in an attempt to force the computer to work the way it did previously. Restore points can be created manually and are also created automatically by the operating system before new applications, updates, or hardware are installed. macOS uses the Time Machine utility, which works in a similar manner. Though there is no similar tool in Linux, you can back up the ~/home directory to a separate partition. When these contents are decompressed to a new install, most of the Linux system and settings are restored. Another option in general is to use imaging (cloning) software. Remember that these techniques do not necessarily back up data and that you should treat the data as a separate entity that needs to be backed up regularly.

- **Consider whole disk encryption:** Finally, you can use whole disk encryption to secure the contents of the drive, making it harder for attackers to obtain and interpret its contents.

A recommendation I give to all my students and readers is to separate the operating system from the data physically. If you can have each on a separate hard drive, it can make things a bit easier just in case the OS is infected with malware (or otherwise fails). The hard drive that the OS inhabits can be completely wiped and reinstalled without worrying about data loss, and applications can always be reloaded. Of course, you should back up settings (or store them on the second drive). If a second drive isn't available, consider configuring the one hard drive as two partitions—one for the OS (or system) and one for the data. By doing this and keeping a well-maintained computer, you are effectively hardening the OS.

By maintaining the workstation or server, you are hardening it as well. This process can be broken down into six steps (and one optional step):

Key Topic

Step 1. **Use a surge protector or UPS.** Make sure the computer and other equipment connect to a surge protector, or better yet a UPS if you are concerned about power loss.

Step 2. **Update the BIOS and/or UEFI.** Flashing the BIOS isn't always necessary. Check the manufacturer's website for your motherboard to see if an update is needed.

Step 3. **Update the OS.** For Windows, this step includes any combination hotfixes and any Windows Updates beyond that, and setting Windows to alert if there are any new updates. For Linux and macOS, it means simply updating the system to the latest version and installing individual patches as necessary.

Step 4. **Update antimalware.** This step includes making sure that there is a current license for the antimalware (antivirus and antispyware) and verifying that updates are turned on and the software is regularly scanning the system.

Step 5. **Update the firewall.** Be sure to have some kind of firewall installed and enabled; then update it. If it is Windows Defender Firewall, updates should happen automatically through Windows Updates. However, if you have a SOHO router with a built-in firewall or other firewall device, you need to update the device's ROM by downloading the latest image from the manufacturer's website.

Step 6. **Maintain the disks.** This instruction means running a disk cleanup program regularly and checking to see whether the hard disk needs to be defragmented from once a week to once a month depending on the amount of usage. It also means creating restore points, doing computer backups, or using third-party backup or drive imaging software.

Step 7. **(Optional) Create an image of the system.** After all your configurations and hardening of the OS are complete, you might consider creating an image of the system. Imaging the system is like taking a snapshot of the entire system partition. That information is saved as one large file, or a set of compressed files that can be saved anywhere. It's kind of like system restore but at another level. The beauty of this process is that you can reinstall the entire image if your system fails or is compromised, quickly and efficiently, with very little configuration necessary. You need to apply only the latest security and AV updates since the image was created. Of course, most imaging software has a price tag involved, but the investment can be well worth the cost if you are concerned about the time needed to get your system back up and running in the event of a failure. This is the basis for standardized images in many organizations. By applying mandated security configurations, updates, and so on, and then taking an image of the system, you can create a snapshot in time that you can easily revert to if necessary, while being confident that a certain level of security is already embedded into the image.

NOTE To clean out a system regularly, consider reimaging it, or if it's a mobile device, resetting it. Reimaging takes care of any pesky malware by deleting everything and reinstalling to the point in time of the image, or to factory condition. Although you will have to do some reconfigurations, the system will also run much faster because it has been completely cleaned out.

Improper or Weak Patch Management

Key Topic

Patch management is the planning, testing, implementing, and auditing of patches. You must adopt a proper patch management program to include all your *operating systems*, *firmware* of your devices in your infrastructure, and *applications* (on-premises and in the cloud).

To be considered secure, operating systems should have support for multilevel security and be able to meet government requirements. An operating system that meets these criteria is known as a ***trusted operating system (TOS)***. Examples of certified trusted operating systems include Windows 10, SELinux, FreeBSD (with the Trusted-BSD extensions), and Red Hat Enterprise Server. For an OS to be considered a TOS, the manufacturer of the system must have strong policies concerning updates and patching.

Even without being a TOS, operating systems should be updated regularly. For example, Microsoft recognizes the deficiencies in an OS and possible exploits that

could occur and therefore releases patches to increase OS performance and protect the system.

Before you update, the first thing to do is to find out the version number, build number, and the patch level. For example, in Windows you can find out this information by opening the System Information tool (open the Run prompt and type **msinfo32.exe**). It is listed directly in the System Summary. You could also use the **winver** command. In addition, you can use the **systeminfo** command at the Command Prompt (a *great* information gatherer!) or simply the **ver** command.

Then you should check your organization's policies to see what level a system should be updated to. You should also test updates and patches on systems located on a clean, dedicated, testing network before going live with an update.

Windows uses the Windows Update program to manage updates to the system. Versions before Windows 10 include this feature in the Control Panel. It can also be accessed by typing **wuapp.exe** at the Run prompt. In Windows 10 it is located in the Settings section or can be opened by typing **ms-settings:windowsupdate** at the Run prompt.

There are various options for the installation of Windows updates, including Automatic Install, Defer Updates to a Future Date, Download with Option to Install, Manual Check for Updates, and Never Check for Updates. The types available to you will differ depending on the version of Windows. Businesses usually frown on Automatic Install, especially in the enterprise. Generally, as the security administrator, you want to specify what is updated, when it is updated, and to which computers it is updated. This way you eliminate problems such as patch level mismatch and network usage issues.

In some cases, your organization might opt to turn off Windows Update altogether. Depending on your version of Windows, turning it off may or may not be possible within the Windows Update program. However, your organization might go a step further and specify that the Windows Update *service* be stopped and disabled, thereby disabling updates. This can be done in the Services console window (Run > **services.msc**) or from the Command Prompt with one of the methods discussed earlier in the chapter; the service name for Windows Update is wuauserv.

Patches and Hotfixes

The best place to obtain patches and hotfixes is from the manufacturer's website. The terms *patches* and *hotfixes* are often used interchangeably. Windows updates are made up of hotfixes. Originally, a *hotfix* was defined as a single problem-fixing patch to an individual OS or application installed live while the system was up and running and without needing a reboot. However, the meaning of this term has changed over time and varies from vendor to vendor. (Vendors may even use both terms to

describe the same thing.) For example, if you run the **systeminfo** command at the Command Prompt of a Windows computer, you see a list of hotfixes. They can be identified with the letters *KB* followed by seven numbers. Hotfixes can be single patches to individual applications, or they might affect the entire system.

On the other side of the spectrum, a gaming company might define hotfixes as a "hot" change to the server with no downtime, and no client download is necessary. The organization releases them if they are critical, instead of waiting for a full patch version. The gaming world commonly uses the terms *patch version*, *point release*, or *maintenance release* to describe a group of file updates to a particular gaming version. For example, a game might start at version 1 and later release an update known as 1.17. The .17 is the point release. (This could be any number, depending on the number of code rewrites.) Later, the game might release 1.32, in which .32 is the point release, again otherwise referred to as the patch version. This naming convention is common with other programs as well.

Administrators should keep up with software and hardware provider security advisories and disclosed vulnerabilities. This keeps you "in the know" when it comes to the latest updates. And this advice applies to server and client operating systems, server add-ons such as Microsoft Exchange or SQL Server, Office programs, web browsers, and the plethora of third-party programs that an organization might use. Your job just got a bit busier!

Of course, you are usually not concerned with updating games in the working world; you should remove them from a computer if they are found (unless perhaps you work for a gaming company). Multimedia software such as Camtasia, however, is prevalent in some companies, and web-based software such as a bulletin-board system is also common and susceptible to attack.

Patches generally carry the connotation of a small fix in the mind of the user or system administrator, so larger patches are often referred to as software updates, service packs, or something similar. However, if you were asked to fix a single security issue on a computer, a patch would be the solution you would want. For example, various Trojans attack older versions of Microsoft Office for Mac. To counter them, Microsoft released a specific patch for those versions of Office for Mac that disallow remote access by the Trojans.

Before installing an individual patch, you should determine whether it perhaps was already installed as part of a group update. For example, you might read that a particular version of macOS had a patch released for iTunes, and being an enthusiastic iTunes user, you might consider installing the patch. But you should first find out the version of the OS you are running; it might already include the patch. To find this information, simply click the Apple menu and then click About This Mac.

Remember: All systems need to be patched at some point. It doesn't matter if they are Windows, macOS, Linux, UNIX, Android, iOS, hardware appliances, kiosks, automotive computers…need I go on?

Unfortunately, sometimes patches are designed poorly, and although they might fix one problem, they could possibly create another, which is a form of software regression. Because you never know exactly what a patch to a system might do, or how it might react or interact with other systems, it is wise to incorporate patch management.

Patch Management

It is not wise to go running around the network randomly updating computers…not to say that you would do so, however! Patching, like any other process, should be managed properly. The process of patch management includes the planning, testing, implementing, and auditing of patches. I use the following four steps; other companies might have a slightly different patch management strategy, but each of the four concepts should be included:

Key Topic

- **Planning:** Before actually doing anything, you should set a plan into motion. The first thing that needs to be decided is whether the patch is necessary and whether it is compatible with other systems. The Microsoft Security Compliance Toolkit (SCT) is one example of a program that can identify security misconfigurations on the computers in your network, letting you know whether patching is needed. If the patch is deemed necessary, the plan should consist of a way to test the patch in a "clean" network on clean systems, how and when the patch will be implemented, and how the patch will be checked after it is installed.
- **Testing:** Before you automate the deployment of a patch among a thousand computers, it makes sense to test it on a single system or small group of systems first. These systems should be reserved for testing purposes only and should not be used by "civilians" or regular users on the network. I know, this is asking a lot, especially given the number of resources some companies have. But the more you can push for at least a single testing system that is not a part of the main network, the less you will be to blame if a failure occurs!
- **Implementing:** If the test is successful, the patch should be deployed to all the necessary systems. In many cases larger updates are done in the evening or over the weekend. Patches can be deployed automatically using software such as Microsoft's Endpoint Configuration Manager and third-party patch management tools.

- **Auditing:** When the implementation is complete, the systems (or at least a sample of systems) should be audited—first, to make sure the patch has taken hold properly, and second, to check for any changes or failures due to the patch. Microsoft's Endpoint Configuration Manager and third-party tools can be used in this endeavor.

NOTE The concept of patch management, in combination with other application/OS hardening techniques, is collectively referred to as *configuration management*.

Some Linux-based and Mac-based programs and services also were developed to help manage patching and the auditing of patches. Red Hat has services to help system administrators with all the RPMs they need to download and install, which can become a mountain of work quickly! And for those people who run GPL Linux, third-party services are also available. A network with a lot of mobile devices benefits greatly from the use of an MDM platform. Even with all these tools at an organization's disposal, sometimes patch management is just too much for one person, or for an entire IT department, and an organization might opt to contract that work out.

Several standards such as the *Common Security Advisory Framework (CSAF)* and the *Open Vulnerability and Assessment Language (OVAL)* are designed to provide a machine-readable format of security advisories to allow automated assessment of vulnerabilities. CSAF is also the name of the technical committee in the OASIS standards organization. CSAF is the successor of the *Common Vulnerability Reporting Framework (CVRF)*. CSAF enables different stakeholders across different organizations to share critical security-related information in a single format, speeding up information exchange and digestion.

OVAL is an international standard maintained by the Center for Internet Security (CIS). OVAL was originally created by MITRE and then transferred to CIS. You can obtain detailed information about the OVAL specification at https://oval.cisecurity.org. In short, OVAL can be defined in two parts: the OVAL Language and the OVAL Interpreter. OVAL is not a language like C++ but is an XML schema that defines and describes the XML documents to be created for use with OVAL.

OVAL has several uses, one of which is as a tool to standardize security advisory distributions. Software vendors need to publish vulnerabilities in a standard, machine-readable format. By including an authoring tool, definitions repository, and definition evaluator, OVAL enables users to regulate their security advisories. Other uses for OVAL include vulnerability assessment, patch management, auditing, threat indicators, and so on.

Key Topic

Legacy Platforms

You might be surprised at the number of small, medium, and large organizations that are still using ***legacy platforms*** and devices that have passed the vendor's last day of software and hardware support. These legacy devices often are core infrastructure devices such as routers and switches. If you run devices that have passed the last day that a vendor will provide software and hardware fixes, it is almost guaranteed that you will be running vulnerable devices because when devices are past the last day of support, vendors will not investigate or patch security vulnerabilities in those devices. This, in turn, introduces a huge risk to companies and service providers that run these devices.

The Impact of Cybersecurity Attacks and Breaches

It is obvious that major cybersecurity attacks and breaches can be catastrophic to many organizations. You just have to log in to your favorite news site or even Twitter to see the millions of dollars lost because of cybersecurity incidents and breaches.

The following are some of the most common categories of the impact of cybersecurity attacks and breaches:

Key Topic

- **Data loss, data breaches, and data exfiltration:** Data loss caused by a breach can range from just one record to millions of records stolen by an attacker. Attackers can exfiltrate data from an organization using different techniques. They also can perform different types of obfuscation and evasion techniques to go undetected (including encoding of data, tunneling, and encryption). The impact on privacy because of a data breach can be very significant. Some of the data from a data breach can be "replaced" (such as a credit card number), but other data (like a Social Security number or health record) cannot.
- **Identity theft:** Criminals often use breached data for identity theft. They use or buy personal records and data from illegal sites in the dark web and other sources to steal the identity of individuals.
- **Financial:** The impact and consequences of a breach or a cybersecurity incident can lead to fines and lawsuits against the company. In some cases, even executives could be fined or sued.
- **Reputation:** The brand and reputation of a company can also be damaged by major cybersecurity incidents and breaches. Customers can lose confidence and trust in the company that has lost their records because of a cybersecurity breach and in some cases because of negligence.

- **Availability loss:** Cybersecurity incidents can also lead to denial-of-service conditions that, in turn, can also cause significant financial impact. Outages from cybersecurity incidents could be catastrophic for companies and, in some cases, their customers.

When dealing with dollars, risk assessments should be based on a quantitative measurement of risk, impact, and asset value. This is why you have to build a good impact assessment methodology when you perform risk analysis. An excellent tool to create during risk assessment is a risk register, also known as a risk log, which helps to track issues and address problems as they occur. After the initial risk assessment, you will continue to use and refer to the risk register. It can be a great tool for just about any organization but can be of more value to certain types of organizations, such as manufacturers that utilize a supply chain. In this case, the organization would want to implement a specialized type of risk management called *supply chain risk management (SCRM)*. In this approach, the organization collaborates with suppliers and distributors to analyze and reduce risk.

When it comes to risk assessment, qualitative risk analysis is an assessment that assigns numeric values to the probability of a risk and the impact it can have on the system or network. Unlike its counterpart, quantitative risk assessment, it does not assign monetary values to assets or possible losses. It is the easier, quicker, and cheaper way to assess risk but cannot assign asset value or give a total for possible monetary loss. You will learn more about risk assessment in Chapter 34, "Summarizing Risk Management Processes and Concepts."

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 6-4 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 6-4 Key Topics for Chapter 6

Key Topic Element	Description	Page Number
Paragraph	Defining cloud computing	138
List	Listing the cloud computing service models	138
Table 6-2	Programming Vulnerabilities and Attacks	149

Key Topic Element	Description	Page Number
List	Listing the most prevalent types of weak configurations	150
Table 6-3	Ports and Their Associated Protocols	152
Paragraph	Description of the role of vendor management and system management	155
Paragraph	Description of outsourced code development and security concerns	155
Paragraph	Description of the lack of vendor support that puts an organization at risk	156
Paragraph	Description of data storage concerns	156
List	Listing several strategies when maintaining and hardening a hard disk	157
List	The process of maintaining the workstations and servers	159
Paragraph	Description of patch management	160
List	Listing the patch management strategy phases	163
Section	Legacy Platforms	165
List	Surveying some of the most common categories of the impact of cybersecurity attacks and breaches	165

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

on-premises, cloud computing, software as a service (SaaS), infrastructure as a service (IaaS), platform as a service (PaaS), public cloud, private cloud, hybrid cloud, community cloud, cloud access security broker (CASB), Common Vulnerabilities and Exposures (CVE), zero-day vulnerability, vendor management, system integration, outsourced code development, lack of vendor support, supply chain, data storage, patch management, Trusted Operating System (TOS), patches, legacy platforms

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What software tool or service acts as the gatekeeper between a cloud offering and the on-premises network, allowing an organization to extend the reach of its security policies beyond its internal infrastructure?
2. What is the name given to a type of vulnerability that is disclosed by an individual or exploited by an attacker before the creator of the software can create a patch to fix the underlying issue?
3. What cloud architecture model is a mix of public and private, but one where multiple organizations can share the public portion?
4. What type of vulnerability occurs when an attacker obtains control of a target computer through some sort of vulnerability, gaining the power to execute commands on that remote computer?
5. What protocol uses TCP ports 465 or 587 in most cases?

This page intentionally left blank



This chapter covers the following topics related to Objective 1.7 (Summarize the techniques used in security assessments) of the CompTIA Security+ SY0-601 certification exam:

- Threat hunting
  - Intelligence fusion
  - Threat feeds
  - Advisories and bulletins
  - Maneuver
- Vulnerability scans
  - False positives
  - False negatives
  - Log reviews
  - Credentialed vs. non-credentialed
  - Intrusive vs. non-intrusive
  - Application
  - Web application
  - Network
  - Common Vulnerabilities and Exposures (CVE) and Common Vulnerability Scoring System (CVSS)
  - Configuration review
- Syslog/Security information and event management (SIEM)
  - Review reports
  - Packet capture
  - Data inputs
  - User behavior analysis
  - Sentiment analysis
  - Security monitoring
  - Log aggregation
  - Log collectors
- Security orchestration, automation, and response (SOAR)

CHAPTER 7

Summarizing the Techniques Used in Security Assessments

This chapter starts by introducing threat hunting and how the threat-hunting process leverages threat intelligence. Then you learn about vulnerability management tasks, such as keeping up with security advisories and performing vulnerability scans. You also learn about the importance of collecting logs (such as system logs [syslogs]) and analyzing those logs in a Security Information and Event Management (SIEM) system. In addition, you learn how security tools and solutions have evolved to provide Security Orchestration, Automation, and Response (SOAR) capabilities to better defend your network, your users, and your organizations overall.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 7-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 7-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Threat Hunting	1–3
Vulnerability Scans	4–6
Syslog and Security Information and Event Management (SIEM)	7–8
Security Orchestration, Automation, and Response (SOAR)	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What is the act of proactively and iteratively looking for threats in your organization that may have bypassed your security controls and monitoring capabilities?
   a. Threat intelligence
   b. Threat hunting
   c. Threat binding
   d. None of these answers are correct.

2. Which of the following provides a matrix of adversary tactics, techniques, and procedures that modern attackers use?
   a. ATT&CK
   b. CVSS
   c. CVE
   d. All of these answers are correct.

3. Which identifier is assigned to disclosed vulnerabilities?
   a. CVE
   b. CVSS
   c. ATT&CK
   d. TTP

4. Which broad term describes a situation in which a security device triggers an alarm, but no malicious activity or actual attack is taking place?
   a. False negative
   b. True negative
   c. False positive
   d. True positive

5. Which of the following is a successful identification of a security attack or a malicious event?
   a. True positive
   b. True negative
   c. False positive
   d. False negative

6. Which of the following occurs when a vulnerability scanner logs in to the targeted system to perform deep analysis of the operating system, running applications, and security misconfigurations?
   a. Credentialed scan
   b. Application scan
   c. Noncredentialed scan
   d. None of these answers are correct.

7. Which of the following are functions of a SIEM?
   a. Log collection
   b. Log normalization
   c. Log correlation
   d. All of these answers are correct.

8. Which solution allows security analysts to collect network traffic metadata?
   a. NetFlow
   b. SIEM
   c. SOAR
   d. None of these answers are correct.

9. Which solution provides capabilities that extend beyond traditional SIEMs?
   a. SOAR
   b. CVSS
   c. CVE
   d. IPFIX

10. Which of the following can be capabilities and benefits of a SOAR solution?
    a. Automated vulnerability assessment
    b. SOC playbooks and runbook automation
    c. Orchestration of multiple SOC tools
    d. All of these answers are correct.

Foundation Topics

Threat Hunting

No security product or technology in the world can detect and block all security threats in the continuously evolving threat landscape (regardless of the vendor or how expensive it is). This is why many organizations are tasking senior analysts in their computer security incident response team (CSIRT) and their security operations center (SOC) to hunt for threats that may have bypassed any security controls that are in place. This is why threat hunting exists.

Key Topic

Threat hunting is the act of proactively and iteratively looking for threats in your organization. This chapter covers details about threat-hunting practices, the operational challenges of a threat-hunting program, and the benefits of a threat-hunting program.

The threat-hunting process requires deep knowledge of the network and often is performed by SOC analysts (otherwise known as investigators, threat hunters, tier 2 or tier 3 analysts, and so on). Figure 7-1 illustrates the traditional SOC tiers and where threat hunters typically reside. In some organizations (especially small organizations), threat hunting could be done by anyone in the SOC because the organization may not have a lot of resources (analysts). The success of threat hunting completely depends on the maturity of the organization and the resources available.

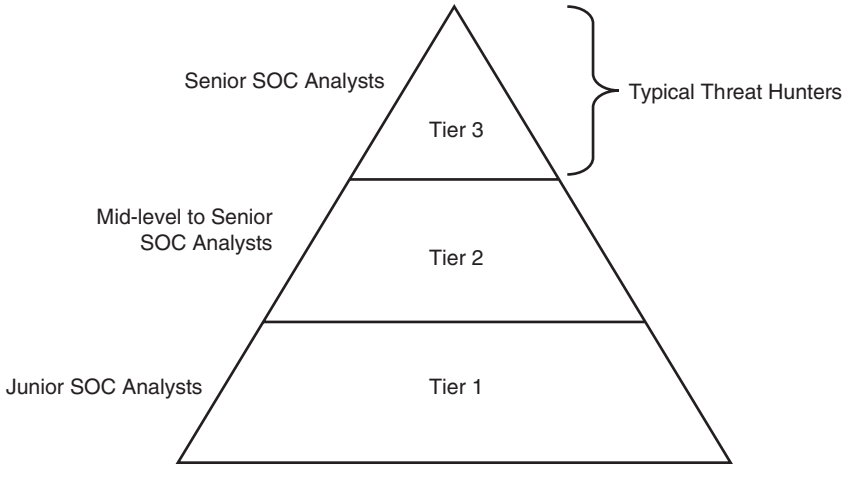


FIGURE 7-1 The SOC Tiers

Some organizations might have a dedicated team within or outside the SOC to perform threat hunting. However, one of the common practices is to have the hunters embedded within the SOC.

Threat hunters assume that an attacker has already compromised the network. Consequently, they need to come up with a hypothesis of what is compromised and how an adversary could have performed the attack. For the threat hunting to be successful, hunters need to be aware of the adversary tactics, techniques, and procedures (TTPs) that modern attackers use. This is why many organizations use MITRE's ATT&CK framework to be able to learn about the tactics and techniques of adversaries. Later in this chapter you learn more about how MITRE's ATT&CK can be used in threat hunting.

Threat hunting is not a new concept. Many organizations have performed threat hunting for a long time. However, in the last decade many organizations have adopted new ways to enhance the threat-hunting process with automation and orchestration.

Threat hunting is not the same as the traditional SOC incident response (reactive) activities. Threat hunting is also not the same as vulnerability management (the process of patching vulnerabilities across the systems and network of your organization, including cloud-based applications in some cases). However, some of the same tools and capabilities may be shared among threat hunters, SOC analysts, and vulnerability management teams. Tools and other capabilities such as data analytics, TTPs, vulnerability feeds, and ***threat feeds*** may be used across the different teams and analysts in an organization.

A high-level threat-hunting process includes the following steps:

Step 1. Threat hunting starts with a trigger based on an anomaly, threat intelligence, or a hypothesis (what could an attacker have done to the organization?). From that moment you should ask yourself: "Do we really need to perform this threat-hunting activity?" or "What is the scope?"

Step 2. Then you identify the necessary tools and methodologies to conduct the hunt.

Step 3. Once the tools and methodologies are identified, you reveal new attack patterns, TTPs, and so on.

Step 4. You refine your hunting tactics and enrich them using data analytics. Steps 2–3 can take one cycle or be iterative and involve multiple loops (depending on what you find and what additional data and research need to be done).

Step 5. A successful outcome could be that you identify and mitigate the threat. However, you need to recognize that in some cases this may not be the case. You may not have the necessary tools and capabilities, or there was no actual threat. This is why the success of your hunting program depends on the maturity of your capabilities and organization as a whole.

You can measure the maturity of your threat-hunting program within your organization in many ways. Figure 7-2 shows a matrix that can be used to evaluate the maturity level of your organization against different high-level threat-hunting elements.

These threat-hunting maturity levels can be categorized as easily as level 1, 2, and 3, or more complex measures can be used.

When it comes to threat intelligence and threat hunting, automation is key! Many organizations are trying to create threat ***intelligence fusion*** techniques to automatically extract threat intelligence data from heterogeneous sources to analyze such data. The goal is for the threat hunter and network defender to maneuver quickly—and faster than the attacker. This way, you can stay one step ahead of threat actors and be able to mitigate the attack.

Security Advisories and Bulletins

Key Topic

In Chapter 5, "Understanding Different Threat Actors, Vectors, and Intelligence Sources," you learned how vendors, coordination centers, security researchers, and others publish ***security advisories*** and bulletins to disclose vulnerabilities. Most of the vulnerabilities disclosed to the public are assigned ***Common Vulnerability and Exposure (CVE)*** identifiers. CVE is a standard created by MITRE (www.mitre.org) that provides a mechanism to assign an identifier to vulnerabilities so that you can correlate the reports of those vulnerabilities among sites, tools, and feeds.

NOTE You can obtain additional information about CVE at https://cve.mitre.org.

One of the most comprehensive and widely used vulnerability databases is the National Vulnerability Database (NVD) maintained by the National Institute of Standards and Technology (NIST). NVD provides information about vulnerabilities disclosed worldwide.

NOTE You can access the NVD and the respective vulnerability feeds at https://nvd.nist.gov.

Threat Hunting High-level Elements	Threat Hunting Maturity Level: Initial (Minimal) Level 1	Intermediate Level 2	Innovative and Leading Level 3
Threat Intelligence and Data Collection	Limited access of threat intelligence and collection of data	High collection of certain types of threat intelligence and data	High collection of many types of threat intelligence and data
Hypothesis Creation	Responds only to existing SIEM, IPS/IDS, firewall logs, etc.	Combines traditional logs with TTPs and threat intelligence	Combines traditional logs with TTPs and threat intelligence and develops automated threat risk scoring
Tools and Techniques for Hunting Hypothesis Testing	Reactive alerts and SIEM searches	Simple tools and analytics leveraging some visualizations, but mostly a manual effort	Advanced search capabilities, visualizations, creating new tools and not depending on traditional tools
TTP Detection	None, only traditional SIEM reactive detection	Identification of indicators of compromise (IoCs) and new attack trends	Able to detect adversary TTPs, IoCs, and create automation for the SOC to routinely detect them in the future
Analytics and Automation	None	Limited analytics and automation	Create automated tools for the SOC to routinely detect threats in the future

FIGURE 7-2 Threat-Hunting Maturity

Most mature vendors such as Microsoft, Intel, and Cisco publish security advisories and bulletins in their websites and are CVE Numbering Authorities (CNAs). CNAs can assign CVEs to disclosed vulnerabilities and submit the information to MITRE and subsequently to NVD.

NOTE You can find additional information about CNAs at https://cve.mitre.org/cve/cna.html.

The following links include examples of security advisories and bulletins published by different vendors:

- **Cisco:** https://www.cisco.com/go/psirt
- **Microsoft:** https://www.microsoft.com/en-us/msrc
- **Red Hat:** https://access.redhat.com/security/security-updates
- **Palo Alto:** https://security.paloaltonetworks.com

Vulnerability disclosures in security advisories are often coordinated among multiple vendors. Most of the products and applications developed nowadays use open-source software. Vulnerabilities in open-source software could affect hundreds or thousands of products and applications in the industry. In addition, vulnerabilities in protocols such as TLS, TCP, BGP, OSPF, and WPA could also affect numerous products and software. Patching open-source and protocol-related vulnerabilities among upstream and downstream vendors is not an easy task and requires good coordination. Figure 7-3 shows the high-level process of a coordinated vulnerability disclosure and underlying patching.

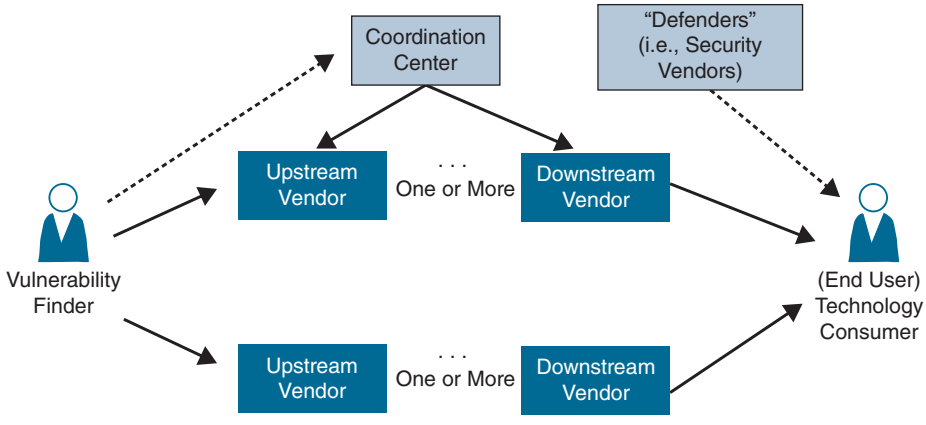


FIGURE 7-3 Coordinated Vulnerability Disclosures

The following steps are illustrated in Figure 7-3:

1. The finder (this can be anyone—a security researcher, customer, security company, an internal employee of a vendor) finds a security vulnerability and reports it to a vendor. The finder can also contact a vulnerability coordination center (such as www.cert.org) to help with the coordination and disclosure.
2. The upstream vendors triage and patch the vulnerability.
3. There could be one or more downstream vendors that also need to patch the vulnerability. In some cases, the coordination center may also interact with downstream vendors in the notification.
4. Security vendors (such as antivirus/antimalware, intrusion detection, and prevention technology providers) may obtain information about the vulnerability and create signatures or any other capabilities to help the end user detect and mitigate an attack caused by the vulnerability.
5. The end user is notified of the patch and the vulnerability.

TIP The preceding process can take days, weeks, months, or even years! Although this process looks very simple in an illustration like the one in Figure 7-3, it is very complicated in practice. For this reason, the Forum of Incident Response and Security Teams (FIRST) has created a Multi-Party Coordination and Disclosure special interest group (SIG) to help address these challenges. You can obtain details about guidelines and practices for multiparty vulnerability coordination and disclosure at https://www.first.org/global/sigs/vulnerability-coordination/multiparty/.

Vulnerability Scans

Vulnerability management teams often use other tools such as vulnerability scanners and software composition analysis (SCA) tools. Figure 7-4 illustrates how a typical automated vulnerability scanner works.

The following are the steps illustrated in Figure 7-4. Keep in mind that vulnerability scanners are all different, but most follow a process like this:

1. In the discovery phase, the scanner uses a tool such as Nmap to perform host and port enumeration. Using the results of the host and port enumeration, the scanner begins to probe open ports for more information.
2. When the scanner has enough information about the open port to determine what software and version are running on that port, it records that information in a database for further analysis. The scanner can use various methods to make this determination, including banner information.

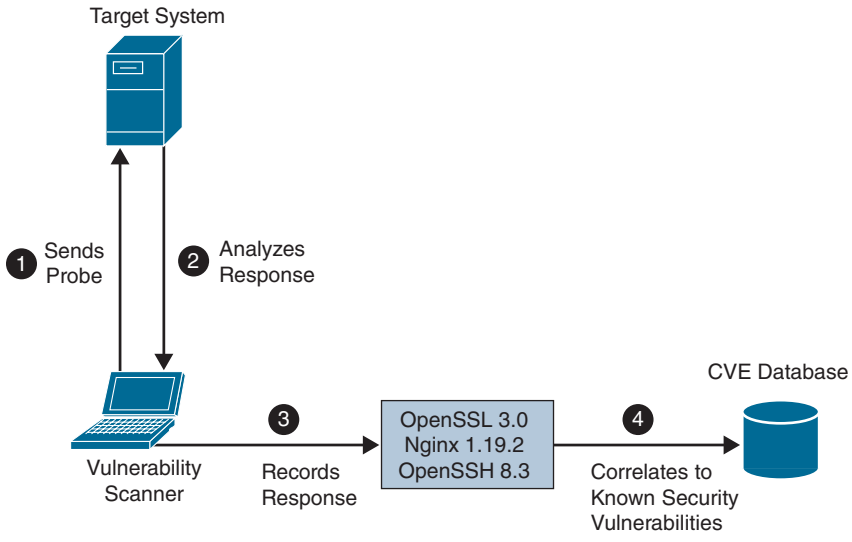


FIGURE 7-4 Coordinated Vulnerability Disclosures

3. The scanner tries to determine if the software that is listening on the target system is susceptible to any known vulnerabilities. It does this by correlating a database of known vulnerabilities against the information recorded in the database about the target services.
4. The scanner produces a report on what it suspects could be vulnerable. Keep in mind that these results are often false positives and need to be validated.

Key Topic

One of the main challenges with automated vulnerability scanners is the number of false positives and false negatives. ***False positive*** is a broad term that describes a situation in which a security device triggers an alarm, but no malicious activity or actual attack is taking place. In other words, false positives are false alarms, and they are also called benign triggers. False positives are problematic because by triggering unjustified alerts, they diminish the value and urgency of real alerts. Having too many false positives to investigate becomes an operational nightmare, and you most definitely will overlook real security events.

There are also ***false negatives***, which is the term used to describe a network intrusion device's inability to detect true security events under certain circumstances—in other words, a malicious activity that is not detected by the security device.

A ***true positive*** is a successful identification of a security attack or a malicious event. A ***true negative*** occurs when the intrusion detection device identifies an activity as acceptable behavior and the activity is actually acceptable.

There are also different types of vulnerability scanners:

- ***Application scanners*:** Used to assess application-specific vulnerabilities and operate at the upper layers of the OSI model
- ***Web application scanners*:** Used to assess web applications and web services (such as APIs)
- ***Network and port scanners*:** Used to determine what TCP or UDP ports are open on the target system

Key Topic

Credentialed vs. Noncredentialed

To reduce the number of false positives, some vulnerability scanners have the capability to log in to a system to perform additional tests and see what programs, applications, and open-source software may be running on a targeted system. These scanners can also ***review logs*** on the target system. They can also perform ***configuration reviews*** to determine if a system may be configured in an unsecure way.

Key Topic

Intrusive vs. Nonintrusive

Vulnerability scanners sometimes can send numerous IP packets at a very fast pace (***intrusive***) to the target system. These IP packets can potentially cause negative effects and even crash the application or system. Some scanners can be configured in such a way that you can throttle the probes and IP packets that it sends to the target system in order to be ***nonintrusive*** and to not cause any negative effects in the system.

Common Vulnerability Scoring System (CVSS)

The Common Vulnerability Scoring System (or ***CVSS***) is an industry standard used to convey information about the severity of vulnerabilities. In CVSS, a vulnerability is evaluated under three aspects, and a score is assigned to each of them. These three aspects (or groups) are the base, temporal, and environmental groups.

- The ***base group*** represents the intrinsic characteristics of a vulnerability that are constant over time and do not depend on a user-specific environment. This is the most important information and the only mandatory information to obtain for a vulnerability score.
- The ***temporal group*** assesses the vulnerability as it changes over time.
- The ***environmental group*** represents the characteristic of a vulnerability taking into account the organization's environment.

The CVSS score is obtained by taking into account the base, temporal, and environmental group information. The score for the base group is between 0 and 10, where 0 is the least severe and 10 is assigned to highly critical vulnerabilities (for example, for vulnerabilities that could allow an attacker to remotely compromise a system and get full control). Additionally, the score comes in the form of a vector string that identifies each of the components used to make up the score. The formula used to obtain the score takes into account various characteristics of the vulnerability and how the attacker is able to leverage these characteristics. CVSS defines several characteristics for the base, temporal, and environmental groups.

TIP You can read and refer to the latest CVSS specification documentation, examples of scored vulnerabilities, and a calculator at www.first.org/cvss.

The base group defines exploitability metrics that measure how the vulnerability can be exploited, and impact metrics that measure the impact on confidentiality, integrity, and availability. In addition to these two, a metric called scope change (S) is used to convey the impact on systems that are affected by the vulnerability but do not contain vulnerable code.

Exploitability metrics include the following:

- **Attack Vector (AV):** Represents the level of access an attacker needs to have to exploit a vulnerability. It can assume four values:
  - Network (N)
  - Adjacent (A)
  - Local (L)
  - Physical (P)
- **Attack Complexity (AC):** Represents the conditions beyond the attacker's control that must exist in order to exploit the vulnerability. The values can be one of the following:
  - Low (L)
  - High (H)
- **Privileges Required (PR):** Represents the level of privileges an attacker must have to exploit the vulnerability. The values are as follows:
  - None (N)
  - Low (L)
  - High (H)

- **User Interaction (UI):** Captures whether user interaction is needed to perform an attack. The values are as follows:
  - None (N)
  - Required (R)
- **Scope (S):** Captures the impact on systems other than the system being scored. The values are as follows:
  - Unchanged (U)
  - Changed (C)

The Impact metrics include the following:

- **Confidentiality Impact (C):** Measures the degree of impact to the confidentiality of the system. It can assume the following values:
  - Low (L)
  - Medium (M)
  - High (H)
- **Integrity Impact (I):** Measures the degree of impact to the integrity of the system. It can assume the following values:
  - Low (L)
  - Medium (M)
  - High (H)
- **Availability Impact (A):** Measures the degree of impact to the availability of the system. It can assume the following values:
  - Low (L)
  - Medium (M)
  - High (H)

The temporal group includes three metrics:

- **Exploit code maturity (E):** Measures whether or not public exploits are available

- **Remediation Level (RL):** Indicates whether a fix or workaround is available
- **Report Confidence (RC):** Indicates the degree of confidence in the existence of the vulnerability

The environmental group includes two main metrics:

- **Security Requirements (CR, IR, AR):** Indicate the importance of confidentiality, integrity, and availability requirements for the system
- **Modified Base Metrics (MAV, MAC, MAPR, MUI, MS, MC, MI, MA):** Allow the organization to tweak the base metrics based on specific characteristics of the environment

For example, a vulnerability that could allow a remote attacker to crash the system by sending crafted IP packets would have the following values for the base metrics:

- Access Vector (AV) would be Network because the attacker can be anywhere and can send packets remotely.
- Attack Complexity (AC) would be Low because it is trivial to generate malformed IP packets.
- Privilege Required (PR) would be None because no privileges are required by the attacker on the target system.
- User Interaction (UI) would also be None because the attacker does not need to interact with any user of the system in order to carry out the attack.
- Scope (S) would be Unchanged if the attack does not cause other systems to fail.
- Confidentiality Impact (C) would be None because the primary impact is on the availability of the system.
- Integrity Impact (I) would be None because the primary impact is on the availability of the system.
- Availability Impact (A) would be High because the device becomes completely unavailable while crashing and reloading.

CVSS also defines a mapping between a CVSS Base Score quantitative value and a qualitative score. Table 7-2 provides the qualitative-to-quantitative score mapping.

Table 7-2 Qualitative-to-Quantitative Score Mapping

Rating	CVSS Base Score
None	0.0
Low	0.1–3.9
Medium	4.0–6.9
High	7.0–8.9
Critical	9.0–10.0

TIP Organizations can use the CVSS score as input to their own risk management processes to evaluate the risk related to a vulnerability and then prioritize the vulnerability remediation.

Logs and Security Information and Event Management (SIEM)

Key Topic

Security Information and Event Management (SIEM) is a specialized device or software used for ***security monitoring***; it collects, correlates, and helps security analysts analyze logs from multiple systems. SIEM typically allows for the following functions:

- **Log collection:** This includes receiving information from devices with multiple protocols and formats, storing the logs, and providing historical reporting and log filtering. A ***log collector*** is software that is able to receive logs from multiple sources (***data input***) and in some cases offers storage capabilities and log analysis functionality.
- **Log normalization:** This function extracts relevant attributes from logs received in different formats and stores them in a common data model or template. This allows for faster event classification and operations. Non-normalized logs are usually kept for archive, historical, and forensic purposes.
- ***Log aggregation*:** This function aggregates information based on common information and reduces duplicates.
- **Log correlation:** This is probably one of the most important SIEM functions. It refers to the capability of the system to associate events gathered by various systems, in different formats and at different times, and create a single actionable event for the security analyst or investigator. Often the quality of SIEM is related to the quality of its correlation engine.

- **Reporting:** Event visibility is also a key functionality of SIEM. Reporting capabilities usually include real-time monitoring and historical base reports.

Most modern SIEMs also integrate with other information systems to gather additional contextual information to feed the correlation engine. For example, they can integrate with an identity management system to get contextual information about users or with NetFlow collectors to get additional flow-based information.

NOTE NetFlow is a technology created by Cisco to collect network metadata about all the different "flows" of traffic on your network. There's also the Internet Protocol Flow Information Export (***IPFIX***), which is a network flow standard led by the Internet Engineering Task Force (IETF). IPFIX was designed to create a common, universal standard of export for flow information from routers, switches, firewalls, and other infrastructure devices. IPFIX defines how flow information should be formatted and transferred from an exporter to a collector. IPFIX is documented in RFC 7011 through RFC 7015 and RFC 5103. Cisco NetFlow Version 9 is the basis and main point of reference for IPFIX. IPFIX changes some of the terminologies of NetFlow, but in essence they are the same principles of NetFlow Version 9.

Several commercial SIEM systems are available. Here's a list of some commercial SIEM solutions:

- Micro Focus ArcSight
- LogRhythm
- IBM QRadar
- Splunk

Figure 7-5 shows how SIEM can collect and process logs from routers, network switches, firewalls, intrusion detection, and other security products that may be in your infrastructure. It can also collect and process logs from applications, antivirus, antimalware, and other host-based security solutions.

Security operation center analysts and security engineers often collect ***packet captures*** during the investigation of a security incident. Packet captures provide the greatest detail about each transaction happening in the network. Full packet capture has been used for digital forensics for many years. However, most malware and attackers use encryption to be able to bypass and obfuscate their transactions. IP packet metadata can still be used to potentially detect an attack and determine the attacker's tactics and techniques.

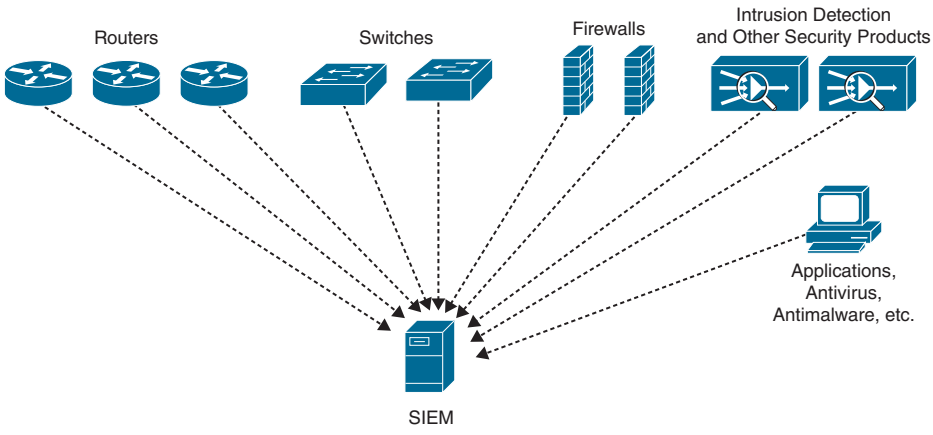


FIGURE 7-5 SIEM Collecting and Processing Logs from Disparate Systems

One of the drawbacks of collecting full packet captures in every corner of your network is the requirement for storage because packet captures in busy networks can take a significant amount of disk space. This is why numerous organizations often collect network metadata with NetFlow or IPFIX and store such data longer than when collecting packet captures.

Several sophisticated security tools also provide ***user behavior analysis*** mechanisms in order to potentially find insiders (internal attackers). Similarly, they provide insights of user behavior even if they do not present a security threat.

Organizations can also deploy ***sentiment analysis*** tools and solutions to help monitor customer sentiment and brand reputation. Often these tools can also reveal the intent and tone behind social media posts, as well as keep track of positive or negative opinions. Threat actors can also try to damage a company's reputation by creating fake accounts and bots in social media platforms like Twitter, Facebook, or Instagram. Attackers can use these fake accounts and bots to provide negative public comments against the targeted organization.

Security Orchestration, Automation, and Response (SOAR)

Key Topic

CSIRT analysts typically work in an SOC utilizing many tools to monitor events from numerous systems (firewalls, applications, IPS, DLP, endpoint security solutions, and so on). Typically, these logs are aggregated in a SIEM. Modern SOCs also use ***Security Orchestration, Automation, and Response (SOAR)*** systems that extend beyond traditional SIEMs.

The tools in the SOC are evolving and so are the methodologies. For example, now security analysts not only respond to basic cyber events but also perform threat hunting in their organizations. SOAR is a set of solutions and integrations designed to allow organizations to collect security threat data and alerts from multiple sources. SOAR platforms take the response capabilities of SIEM to the next level. SOAR solutions supplement, rather than replace, the SIEM. They allow the cybersecurity team to extend its reach by automating the routine work of cybersecurity operations.

TIP Unlike traditional SIEM platforms, SOAR solutions can also be used for threat and vulnerability management, security incident response, and security operations automation.

Deploying SOAR and SIEM together in solutions makes the life of SOC analysts easier. SOAR platforms accelerate incident response detection and eradication times because they can automatically communicate information collected by SIEM with other security tools. Several traditional SIEM vendors are changing their products to offer hybrid SOAR/SIEM functionality.

Another term adopted in the cybersecurity industry is Extended Detection and Response (XDR). XDR is a series of systems working together that collects and correlates data across hosts, mobile devices, servers, cloud workloads, email messages, web content, and networks, enabling visibility and context into advanced threats. The goal of an XDR system is to allow security analysts to analyze, prioritize, hunt, and remediate cybersecurity threats to prevent data loss and security breaches.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 7-3 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 7-3 Key Topics for Chapter 7

Key Topic Element	Description	Page Number
Paragraph	Defining threat hunting	175
Paragraph	Understanding security advisories, bulletins, and what a CVE is	177
Paragraph	Understanding false positives and false negatives	181
Section	Credentialed vs. Noncredentialed	182
Section	Intrusive vs. Nonintrusive	182
Paragraph	Defining what SIEM is	186
Paragraph	Understanding the SOAR concept	188

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

threat hunting, threat feeds, intelligence fusion, security advisories, Common Vulnerability and Exposures (CVE), false positives, false negatives, true positive, true negative, application scanners, web application scanners, network and port scanners, review logs, configuration reviews, intrusive, nonintrusive, CVSS, base group, temporal group, environmental group, Security Information and Event Management (SIEM), security monitoring, log collector, data input, Log aggregation, IPFIX, packet captures, user behavior analysis, sentiment analysis, Security Orchestration, Automation, and Response (SOAR)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What type of vulnerability scanner can be used to assess vulnerable web services?
2. What documents do vendors, vulnerability coordination centers, and security researchers publish to disclose security vulnerabilities?
3. What term is used to describe an organization that can assign CVEs to vulnerabilities?

4. What public database can anyone use to obtain information about security vulnerabilities affecting software and hardware products?
5. How many score "groups" are supported in CVSS?
6. A vulnerability with a CVSS score of 4.9 is considered a _____ severity vulnerability.
7. What is the process of iteratively looking for threats that may have bypassed your security controls?



This chapter covers the following topics related to Objective 1.8 (Explain the techniques used in penetration testing) of the CompTIA Security+ SY0-601 certification exam:

- Penetration testing
  - Known environment
  - Unknown environment
  - Partially known environment
  - Rules of engagement
  - Lateral movement
  - Privilege escalation
  - Persistence
  - Cleanup
  - Bug bounty
  - Pivoting
- Passive and active reconnaissance
  - Drones
  - War flying
  - War driving
  - Footprinting
  - OSINT
- Exercise types
  - Red-team
  - Blue-team
  - White-team
  - Purple-team

Understanding the Techniques Used in Penetration Testing

Penetration testing (otherwise known as ethical hacking) has been extremely popular in the last several years. A penetration tester is someone who mimics what an attacker can do to an organization when finding and exploiting security vulnerabilities. In this chapter, you learn the details about penetration testing methodologies, rules of engagement, and how security researchers also participate in bug bounties. You also learn how penetration testers (pen testers) perform passive and active reconnaissance and the different types of security exercises.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 8-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 8-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Penetration Testing	1–4
Passive and Active Reconnaissance	5–7
Exercise Types	8–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which term is used to define the practice of mimicking a threat actor by using the same methodologies and tools to find and exploit vulnerabilities with the permission of the system or network owner?
   a. Ethical hacking
   b. Pen testing
   c. Penetration testing
   d. All of these answers are correct.

2. Which of the following is a type of penetration testing where the tester starts out with a significant amount of information about the organization and its infrastructure?
   a. Known environment
   b. Unknown environment
   c. Partially known environment
   d. None of these answers are correct.

3. Which of the following are elements of the penetration pre-engagement phase?
   a. Developing the rules of engagement document
   b. Negotiating contracts
   c. Creating the statement of work (SOW)
   d. All of these answers are correct.

4. Which of the following elements are typically included in the rules of engagement document during a penetration testing?
   a. Testing timeline
   b. Location of the testing
   c. The security controls that could potentially detect or prevent testing
   d. All of these answers are correct.

5. Which term is used when a penetration tester uses public records to perform passive reconnaissance?
   a. OSINT gathering
   b. Scanning
   c. Banner fingerprinting
   d. Shodan

6. Which tool can be used to perform active reconnaissance?
   a. Nmap
   b. Nessus
   c. Nikto
   d. All of these answers are correct.

7. Which term is used to describe when attackers or pen testers fly drones to perform reconnaissance of a location or eavesdrop wireless networks?
   a. War driving
   b. War flying
   c. Wireless flying
   d. None of these answers are correct.

8. Which of the following is an example of a blue team?
   a. CSIRT
   b. Pen testing teams
   c. Offensive security teams
   d. None of these answers are correct.

9. What term is used to describe how an organization integrates the defensive capabilities of a blue team with the adversarial techniques used by the red team?
   a. Advanced red teaming
   b. Adversarial emulation
   c. Purple teaming
   d. None of these answers are correct.

10. Which term is often used to define the team that focuses in security governance, regulatory compliance, and risk management?
    a. White team
    b. Purple team
    c. Red team
    d. Blue team

Foundation Topics

Penetration Testing

A penetration tester is typically also referred to as an ethical hacker. The term *ethical hacker* describes a person who acts as an attacker and evaluates the security posture of a computer network for the purpose of minimizing risk. The term *hacker* has been used in many different ways and has many different definitions. Most people in a computer technology field would consider themselves hackers by the simple fact that they like to tinker. This is obviously not a malicious thing. The key factor here in defining ethical versus nonethical hacking is that the latter involves malicious intent. A security researcher looking for vulnerabilities in products, applications, or web services is considered an ethical hacker if he or she responsibly discloses those vulnerabilities to the vendors or owners of the targeted research. However, the same type of "research" performed by someone who then uses the same vulnerability to gain unauthorized access to a target network/system would be considered a nonethical hacker. We could even go so far as to say that someone who finds a vulnerability and discloses it publicly without working with a vendor is considered a nonethical hacker—because this could lead to the compromise of networks/systems by others who use this information in a malicious way.

The truth is that, as an ethical hacker, you use the same tools to find vulnerabilities and exploit targets as do nonethical hackers. However, as an ethical hacker, you would typically report your findings to the vendor or customer you are helping to make more secure. You would also try to avoid performing any tests or exploits that might be destructive in nature. An ethical hacker's goal is to analyze the security posture of a network's or system's infrastructure in an effort to identify and possibly exploit any security weaknesses found and then determine if a compromise is possible. This process is called security ***penetration testing*** or ethical hacking.

So, why do we need penetration testing? Well, first of all, as someone who is responsible for securing and defending a network/system, you want to find any possible paths of compromise before the bad guys do. For years we have developed and implemented many different defensive techniques (for instance, antivirus programs, firewalls, intrusion prevention systems [IPSs], antimalware). We have deployed defense-in-depth as a method to secure and defend our networks. But how do we know if those defenses really work and whether they are enough to keep out the bad guys? How valuable is the data that we are protecting, and are we protecting the right things? These are some of the questions that should be answered by a penetration test.

If you build a fence around your yard with the intent of keeping your dog from getting out, maybe it only needs to be 4 feet tall. However, if your concern is not the dog getting out but an intruder getting in, then you need a different fence—one that would need to be much taller than 4 feet. Depending on what you are protecting, you might also want razor wire on the top of the fence to deter the bad guys even more.

When it comes to information security, you need to do the same type of assessments on your networks and systems. You need to determine what you are protecting and whether your defenses can hold up to the threats that are imposed on them. This is where penetration testing comes in. Simply implementing a firewall, an IPS, antimalware, a virtual private network, a web application firewall (WAF), and other modern security defenses isn't enough. You also need to test their validity. And you need to do this on a regular basis. As you know, networks and systems change constantly. This means the attack surface can change as well, and when it does, you need to consider reevaluating the security posture by way of a penetration test.

When talking about penetration testing methods, you are likely to hear the terms ***known environment***, ***unknown environment***, or ***partially known environment*** testing. These terms are used to describe the perspective from which the testing is performed, as well as the amount of information that is provided to the tester:

Key Topic

- **Known environment:** In this environment the tester starts out with a significant amount of information about the organization and its infrastructure. The tester would normally be provided things like network diagrams, IP addresses, configurations, and a set of user credentials. If the scope includes an application assessment, the tester might also be provided the source code of the target application. The idea of this type of test is to identify as many security holes as possible. In an unknown environment test, the scope may be only to identify a path into the organization and stop there. With "known environment" testing, the scope would typically be much broader and include internal network configuration auditing and scanning of desktop computers for defects. Time and money are typically deciding factors in the determination of which type of penetration test to complete. If a company has specific concerns about an application, a server, or a segment of the infrastructure, it can provide information about that specific target to decrease the scope and the amount of time spent on the test but still uncover the desired results. With the sophistication and capabilities of adversaries today, it is likely that most networks will be compromised at some point, and a known environment approach is not a bad option.
- **Unknown environment:** In this environment the tester is typically provided only a very limited amount of information. For instance, the tester may be provided only the domain names and IP addresses that are in scope for a

particular target. The idea of this type of limitation is to have the tester start out with the perspective that an external attacker might have. Typically, an attacker would first determine a target and then begin to gather information about the target, using public information and gaining more and more information to use in attacks. The tester would not have prior knowledge of the target's organization and infrastructure. Another aspect of unknown environment testing is that sometimes the network support personnel of the target may not be given information about exactly when the test is taking place. This type of environment allows for a defense exercise to take place as well, and it also eliminates the issue of a target preparing for the test and not giving a real-world view of how the security posture really looks.

- **Partially known environment:** This type of penetration test is somewhat of a hybrid approach between the preceding two types. The penetration testers may be provided credentials but not full documentation of the network infrastructure. This way, the testers could still provide results of their testing from the perspective of an external attacker's point of view. Considering the fact that most compromises start at the client and work their way throughout the network, a good approach would be a scope where the testers start on the inside of the network and have access to a client machine. Then they could pivot throughout the network to determine what the impact of a compromise would be.

A number of penetration testing methodologies have been around for a while and continue to be updated as new threats emerge. The following are some of the most common:

- **Penetration Testing Execution Standard (PTES):** www.pentest-standard.org/index.php/Main_Page
- **NIST Technical Guide to Information Security Testing and Assessment:** https://csrc.nist.gov/publications/detail/sp/800-115/final
- **OWASP Web Security Testing Guide:** https://owasp.org/www-project-web-security-testing-guide
- **Open Source Security Testing Methodology Manual (OSSTMM):** www.isecom.org/research.html

Figure 8-1 shows the penetration testing lifecycle based on the Penetration Testing Execution Standard guidance.

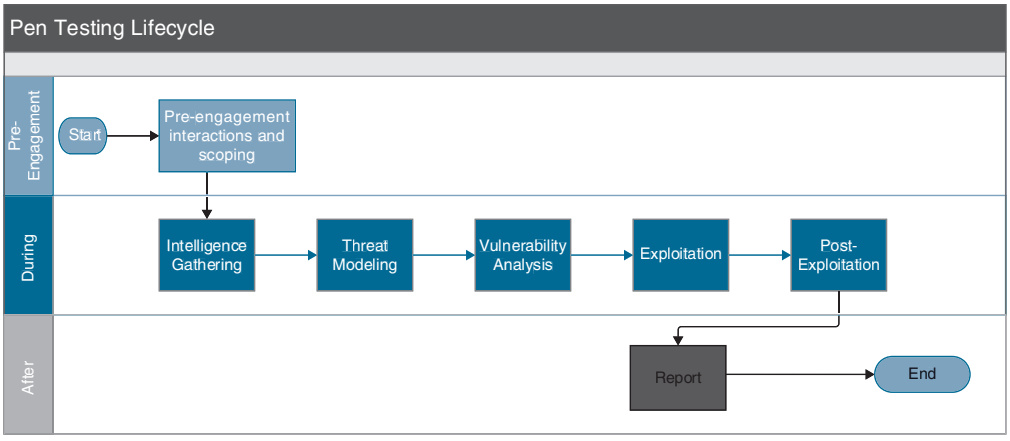


FIGURE 8-1 The Penetration Testing Lifecycle

The penetration testing lifecycle has three major milestones (as shown in Figure 8-1): the pre-engagement tasks, the actual testing, and the report delivered to the customer after testing. The pre-engagement tasks include items such as contract negotiations, the statement of work (SOW), scoping, and the rules of engagement.

Key Topic

The ***rules of engagement*** documentation specifies the conditions under which the security penetration testing engagement will be conducted. You need to document and agree upon these rules of engagement conditions with the client or an appropriate stakeholder. The following elements are typically included in a rules of engagement document:

- Testing timeline
- Location of the testing
- Time window of the testing
- Preferred method of communication
- The security controls that could potentially detect or prevent testing
- IP addresses or networks from which testing will originate
- The scope of the engagement

During the penetration testing, the ethical hacker performs passive and active reconnaissance to find vulnerabilities. Passive and active reconnaissance are discussed later in this chapter. Based on the findings from the reconnaissance phase, the ethical hacker then starts finding security vulnerabilities and demonstrates how an attacker could exploit such vulnerabilities. After exploiting a vulnerability, the ethical hacker

could also demonstrate how an attacker could perform post-exploitation activities such as

Key Topic

- ***Lateral movement* and *Pivoting*:** Lateral movement (also referred to as *pivoting*) is a post-exploitation technique that can be performed using many different methods. The main goal of lateral movement is to move from one device to another to avoid detection, steal sensitive data, and maintain access to these devices to exfiltrate the sensitive data. Lateral movement involves scanning a network for other systems, exploiting vulnerabilities in other systems, compromising credentials, and collecting sensitive information for exfiltration. Lateral movement is possible if an organization does not segment its network in a proper way. Network segmentation is therefore very important. After compromising a system, you can use basic port scans to identify systems or services of interest that you can further attack in an attempt to compromise valuable information.
- ***Privilege escalation*:** This is the process of elevating the level of authority (privileges) of a compromised user or a compromised application. It is done to further perform actions on the affected system or any other systems in the network. It is possible to perform privilege escalation in a few different ways. An attacker may be able to compromise a system by logging in with a non-privileged account. Subsequently, the attacker can go from that unprivileged (or less privileged) account to another account that has greater authority. It is also possible to perform privilege escalation by "upgrading," or elevating, the privileges of the same account.
- ***Persistence*:** After the exploitation phase, you need to maintain a foothold in a compromised system to perform additional tasks such as installing and/or modifying services to connect back to the compromised system. You can maintain persistence of a compromised system in a number of ways, including the following:
  - Creating a bind or reverse shell
  - Creating and manipulating scheduled jobs and tasks
  - Creating custom daemons and processes
  - Creating new users
  - Creating additional backdoors

TIP When you maintain persistence in a compromised system, you can take several actions, such as the following:

- Uploading additional tools
- Using local system tools
- Performing ARP scans and ping sweeps
- Conducting DNS and directory services enumeration
- Launching brute-force attacks
- Performing additional enumeration and system manipulation using management protocols (for example, WinRM, WMI, SMB, SNMP) and compromised credentials
- Executing additional exploits

Key Topic

Ethical hackers or pen testers also need to be able to cover their tracks. Many tools can leave behind residual files or data that you need to be sure to clean from the target systems after the testing phases of a penetration testing engagement are complete. It is also very important to have the client or system owner validate that the ***cleanup*** effort is sufficient. This is not always easy to accomplish, but providing a comprehensive list of activities performed on any systems under test will help.

Following are some items you will want to be sure to clean from systems:

- User accounts created
- Shells spawned on exploited systems
- Database input created by automated tools or manually
- Any tools installed or run from the systems under test

Bug Bounties vs. Penetration Testing

Key Topic

Bug bounties have become very popular in the last several years. In a bug bounty program, an organization provides recognition and compensation to security researchers for reporting security vulnerabilities (which are basically *bugs* in code or hardware). The goal of bug bounty programs is to allow underlying organizations to crowdsource the way they find vulnerabilities in their systems and infrastructure in a scalable way.

These programs allow an underlying organization to learn about bugs and fix them before an attacker can abuse them. Several bug bounty platforms (or brokers) help

crowdsource and manage a bug bounty program. Examples of bug bounty platforms/companies include

- HackerOne (www.hackerone.com)
- BugCrowd (www.bugcrowd.com)
- Intigriti (www.intigriti.com)

Security researchers and ethical hackers often participate either full-time or part-time in these bug bounty programs.

TIP For a list of references and resources about bug bounty programs, platforms, tools, and tips, see https://github.com/The-Art-of-Hacking/h4cker/tree/master/bug-bounties.

Passive and Active Reconnaissance

Reconnaissance can be passive or active. ***Passive reconnaissance*** can be carried out by an attacker just researching information about the victim's public records, social media sites, and other technical information, such as DNS, whois, and sites such as Shodan (www.shodan.io). Searching through these public records is often referred to as ***open-source intelligence (OSINT)*** gathering. The ethical hacker or pen tester can use tools such as Maltego, Recon-ng, theHarvester, SpiderFoot, OWASP Amass, and many others to accelerate this research.

TIP The download links of the aforementioned OSINT tools (along with several others) are included in the GitHub repository at https://github.com/The-Art-of-Hacking/h4cker/tree/master/osint. This GitHub repository includes a curated list of numerous references and resources about OSINT and passive reconnaissance tools and methodologies.

For instance, the Shodan search engine is a powerful database of prescanned networked devices connected to the Internet. It consists of scan results including banners collected from port scans of public IP addresses, with fingerprints of services like Telnet, FTP, HTTP, and other applications.

Shodan creates a risk profile by providing both attackers and defenders with a prescanned inventory of devices connected to public IP addresses on the Internet. For example, when a new vulnerability is discovered and published, an attacker can quickly and easily search Shodan for vulnerable devices and then launch an attack.

Attackers can also search the Shodan database for devices with poor configurations or other weaknesses, all without actively scanning. Using Shodan search filters, a user can really narrow down search results, by country code or classless interdomain routing (CIDR) netblock, for example. Shodan application programming interfaces (APIs) and some basic scripting can enable many search queries and subsequent actions (for example, a weekly query of newly discovered IPs scanned by Shodan on a CIDR netblock that runs automatically and is emailed to the security team). Remember that public IP addresses are constantly probed and scanned already. By using Shodan, you are not scanning because Shodan has already scanned these IPs. Shodan is a tool, and it can be used for good or evil. To mitigate risk, you can take tangible steps like registering for a free Shodan account, searching for your organization's public IPs, and informing the right network and security people of the risks of your organization's Shodan exposure. You can learn more at www.shodan.io.

Active reconnaissance is carried out mostly by using network and vulnerability scanners. The following are commercial and open-source application, port, and vulnerability scanners:

- AppScan by IBM
- Burp Suite Professional by PortSwigger
- Nessus by Tenable Network Security
- Netsparker by Mavituna Security
- Nexpose by Rapid7
- Nmap (open-source port scanner)
- Nikto (open-source web application scanner)
- OWASP Zed Attack Proxy (open-source web application scanner, proxy, and attack platform maintained by the Open Web Application Security Project [OWASP])
- Qualys
- Retina Web Security Scanner by eEye Digital Security
- Sentinel by WhiteHat
- Veracode Web Application Security by Veracode
- VUPEN Web Application Security Scanner by VUPEN Security
- WebApp360 by nCircle

Key Topic

One of the main purposes of reconnaissance is to footprint (perform ***footprinting*** on) an application, system, or network to find vulnerabilities that could potentially be exploited. Attackers and pen testers have used several creative methods to perform reconnaissance. For example, even ***drones*** have been used to eavesdrop and monitor wireless networks. The term ***war driving*** refers to the ability of an attacker to just drive around and get a huge amount of information over a very short period of time. The attacker could obtain observations about wireless access points, underlying service set identifiers (SSIDs), and even Bluetooth and cellular communications. A good example of a repository and application that has been used to store and analyze war driving records is Wigle (https://wigle.net/). Any individual can download the Wigle mobile app and automatically upload information about Wi-Fi networks, cell towers, and Bluetooth devices.

A similar concept is ***war flying***. In war flying an attacker can fly drones or similar devices to obtain information about a wireless network or even collect pictures of facilities in some cases.

TIP Another great resource to learn about real-life adversarial techniques used for passive and active reconnaissance (as well as the complete lifecycle of a cybersecurity attack) is the MITRE ATT&CK framework, which can be accessed at https://attack.mitre.org.

Exercise Types

Several terms have been used to define the different penetration testing or cybersecurity exercise types. The following terms are the most common:

Key Topic

- ***Red team*:** Individuals who are performing adversarial simulation and penetration testing. However, a "true red team engagement" goes beyond the traditional scope of a penetrating testing. For example, red teamers can also demonstrate how an attacker could infiltrate a building and perform advanced social engineering attacks. Red teamers also create exploits and their own tools.
- ***Blue team*:** Defenders of the organizations. Blue teams typically include the computer security incident response team (CSIRT) and information security (InfoSec) teams. The expression "blue teams versus red teams" is adopted from the military, where the red team is typically the "offensive" team and blue is the "defensive" team.
- ***Purple team*:** Individuals who perform different tactics to maximize the effectiveness of red and blue team operations. If you combine the colors red and

blue, you end up with purple. Similarly in security, purple teams integrate the defensive capabilities of a blue team with the adversarial techniques used by the red team. In most cases, the purple team is not a separate team, but a solid dynamic between the blue and read teams.

- ***White team*:** A team that blends all previous colors together. White teams are individuals who are focused on governance, management, risk assessment, and compliance.

TIP In teaming security assessments, the blue team members are the defenders. It is their job to counter the red team and keep them from accomplishing their mission. This team has the additional advantage of measuring and improving alerting and response. The red team members act as the adversaries. They are the ones attacking and trying to remain unnoticed. The white team members are neutral. They are the referees who define the goals and rules and also adjudicate the exercise. The goals of the purple team are to effectively combine the skills and knowledge of both the red and blue teams to achieve maximum effectiveness.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 8-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 8-2 Key Topics for Chapter 8

Key Topic Element	Description	Page Number
List	Understanding known environment, unknown environment, or partially known environment penetration testing	198
Paragraph	Defining rules of engagement	200
List	Defining lateral movement, pivoting, privilege escalation, and persistence	201
Paragraph	Understanding how attackers and penetration testing cover their tracks and perform cleanup	202

Key Topic Element	Description	Page Number
Paragraph	Surveying what bug bounties are	202
Paragraph	Defining footprinting, war driving, and war flying	205
List	Defining red, blue, purple, and white teams	205

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

penetration testing, known environment, unknown environment, partially known environment, rules of engagement, lateral movement, pivoting, privilege escalation, persistence, cleanup, bug bounties, passive reconnaissance, open-source intelligence (OSINT), active reconnaissance, footprinting, drones, war driving, war flying, red team, blue team, purple team, white team

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. You were hired to perform a penetration test against a set of applications. After the exploitation phase, you need to maintain a foothold in a compromised system to perform additional tasks such as installing and/or modifying services to connect back to the compromised system. This process is referred to as _____.
2. What is the process of elevating the level of authority (privileges) of a compromised user or a compromised application?
3. What is the term used to define the type of testing where the penetration testers may be provided credentials but not full documentation of the network infrastructure?
4. What is the term used when an organization provides recognition or compensation to security researchers and ethical hackers who report security vulnerabilities or bugs? Often organizations can use brokers and companies that manage the compensation and communication with the security researchers.
5. OSINT is used in the _____ reconnaissance phase of the penetration testing lifecycle.



This chapter covers the following topics related to Objective 2.1 (Explain the importance of security concepts in an enterprise environment) of the CompTIA Security+ SY0-601 certification exam:

- Configuration management
  - Diagrams
  - Baseline configuration
  - Standard naming conventions
  - Internet protocol (IP) schema
- Data sovereignty
- Data protection
  - Data loss prevention (DLP)
  - Masking
  - Encryption at rest, in transit/motion, and in processing
  - Tokenization
  - Rights management
- Geographical considerations
- Response and recovery controls
- Secure Sockets Layer (SSL)/Transport Layer Security (TLS) inspection
- Hashing
- API considerations
- Site resiliency
  - Hot site
  - Cold site
  - Warm site
- Deception and disruption
  - Honeypots
  - Honeyfiles
  - Honeynets
  - Fake telemetry
  - DNS sinkhole

CHAPTER 9

Understanding the Importance of Security Concepts in an Enterprise Environment

This chapter starts with an overview of best practices for system configuration management and then details different approaches for data sovereignty and protection. You learn about different geographical considerations for data protection, as well as response and recovery controls. This chapter also covers the principles of Secure Sockets Layer (SSL)/Transport Layer Security (TLS) inspection, hashing, and considerations to protect APIs. You also learn about the different methods for site resiliency and techniques for deception and disruptions (including honeypots, honeyfiles, honeynets, fake telemetry, DNS sinkholes, and others).

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 9-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 9-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Configuration Management	1–3
Data Sovereignty and Data Protection	4–7
Site Resiliency	8–9
Deception and Disruption	10–12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is a primary goal of configuration management?
   a. Maintaining computer systems, servers, network infrastructure, and software in a desired, consistent state
   b. Reducing the cost of acquiring computer systems, servers, network infrastructure, and software used for information security
   c. Ensuring that any changes done to the infrastructure do not affect the underlying organization's IT budget
   d. All of these answers are correct.

2. After a minimum desired state of security is defined, _____ should be taken to assess the current security state of computers, servers, network devices, and the network in general.
   a. network diagrams
   b. IPv4 schemas
   c. baselines
   d. None of these answers are correct.

3. Which of the following is a benefit of standard naming conventions in an IT infrastructure?
   a. Appropriate naming conventions are used to avoid conflicts and to be able to correlate data among disparate systems.
   b. Appropriate naming conventions are used to reduce unnecessary spending of IT infrastructure.
   c. Appropriate naming conventions are used to better create IPv6 network schemas and for data sovereignty.
   d. None of these answers are correct.

4. Which of the following are privacy laws or regulations? (Choose two.)
   a. PCI-DSS
   b. CCPA
   c. GDPR
   d. FedRamp

5. Which of the following is a type of software or hardware-based data loss prevention solution?
   a. Endpoint DLP systems
   b. Network DLP systems
   c. Storage DLP systems
   d. All of these answers are correct.

6. You were hired to deploy a system to prevent unauthorized use and transmission of confidential information. What should you prioritize to protect and encrypt?
   a. Data at rest
   b. Data in use
   c. Data in motion
   d. All of the answers are correct.

7. Which of the following are used in digital signatures, in file and message authentication, and as a way to protect and verify the integrity of sensitive data?
   a. Data masking
   b. Tokenization
   c. Hashes
   d. Redaction

8. Which term is used when you have a near duplicate of the original site of the organization that can be up and running within minutes?
   a. Hot site
   b. Warm site
   c. Cluster site
   d. Cold site

9. What do you call a redundant site that has tables, chairs, bathrooms, and possibly some technical setup, but a lot of configuration of computers and data restoration is necessary before the site can be properly utilized?
   a. Hot site
   b. Warm site
   c. Cluster site
   d. Cold site

10. Which term is used to categorize a group of computers used to attract and trap potential adversaries to counteract and analyze an attack?
   a. Honeypot
   b. Honeynet
   c. Honeyfile
   d. None of these answers are correct.

11. A security analyst creates a file called passwords.txt to lure attackers to access it. Which term is used for this technique?
   a. Honeynet
   b. Honeypot
   c. Honeyfarm
   d. Honeyfile

12. In a _____ you configure one or more DNS servers to provide false results to attackers and redirect them to areas in the network where you can observe their tactics and techniques.
   a. DNS sinkhole
   b. DNS tunnel
   c. DNS Zone transfer
   d. None of these answers are correct.

Foundation Topics

Configuration Management

Configuration management is an ongoing process created with the goal of maintaining computer systems, servers, network infrastructure, and software in a desired, consistent state. One of the primary goals of configuration management is to ensure that your infrastructure performs as it's expected to as changes are made over time.

Several key elements are used in the process of configuration management:

Key Topic

- ***Diagrams* and other documentation:** A good configuration management process helps to avoid small or large changes going undocumented. These undocumented changes can lead to poor performance, inconsistencies, or noncompliance and negatively impact business operations and security. When poorly documented changes are made within your infrastructure, they add to instability and downtime. Having good network diagrams and well-written and up-to-date documentation is crucial and allows you to not only troubleshoot problems but also respond quickly to security incidents.
- ***Baseline configuration*:** After a minimum desired state of security is defined, baselines should be taken to assess the current security state of computers, servers, network devices, and the network in general. Baseline configurations should be properly documented and reviewed to include a set of specifications for information systems or configuration items within those systems. Baseline configurations are used by security professionals, along with network and system administrators, as a basis for future deployments, releases, or changes to information systems and applications. A baseline configuration could include information about standard software packages installed on endpoint systems, servers, network infrastructure devices, mobile devices, or applications and infrastructure hosted in the cloud. These baseline configurations should also include current version numbers and patch information on operating systems and applications, and configuration parameters, network topology, and the logical placement of those components within the system architecture. You should always review configuration baselines to make sure that they are still relevant. New baselines should be created as organizational information systems change over time.
- ***Standard naming conventions*:** You should make sure that your organization has appropriate naming conventions for describing IT infrastructure, applications, and users. Appropriate naming conventions are used to avoid conflicts and to be able to correlate data among disparate systems.

- ***Internet protocol (IP) schema*:** Similar to standard naming conventions, having a proper IPv4 or IPv6 schema will help avoid conflicts within your on-premises network or cloud deployments and to be able to correlate data among disparate systems. For example, when using RFC 1918 private IP addresses, you should perform proper planning so that the IP network scheme can be more organized, easier to set up, and easier to troubleshoot than user and network services. Identify what subnets are used for wired and wireless users, as well as virtual private network (VPN) users. Wireless access may require additional subnets for guest access, quarantine of nonsecure devices, IoT devices, and so on. In addition, you may have IP address subnets dedicated for voice over IP (VoIP) devices, separate from printers and IoT devices. Configuration management systems let you consistently define system settings, as well as build and maintain those systems according to those baseline settings.

Data Sovereignty and Data Protection

Key Topic

Data sovereignty refers to any information (data) that has been converted and stored in a digital form. Many laws and regulations have regulated how organizations should handle their customer and employee data. One of the main concerns around data sovereignty is privacy. For example, the General Data Protection Regulation (GDPR) is a regulation in the European Union (EU) and the European Economic Area (EEA) focused on data protection and privacy. Another example is the California Consumer Privacy Act (CCPA). These regulations give consumers the right to know what personal information is being collected by companies, government, and any other organizations. Consumers can also access the personal information that is collected and request that it be deleted, as well as to know whether their personal information is being shared (and if so, with whom). In addition, consumers can opt out of the sale of their personal information.

Key Topic

Another concept that is crucial for ***data protection*** is ***data loss prevention (DLP)*** systems. DLP is a concept that refers to the monitoring of data in processing, data in transit/motion, and data at rest. A DLP system performs content inspection and is designed to prevent unauthorized use of data as well as prevent the leakage of data outside the computer (or network) where it resides. DLP systems can be software- or hardware-based solutions and come in three varieties:

- **Endpoint DLP systems:** These systems run on an individual computer and are usually software-based. They monitor data in processing, such as email communications, and can control what information flows between various users. These systems can also be used to inspect the content of USB-based mass-storage devices or block those devices from being accessed altogether by creating rules within the software.

- **Network DLP systems:** These software- or hardware-based solutions are often installed on the perimeter of the network. They inspect data that is in transit/motion.
- **Storage DLP systems:** These systems are typically installed in data centers or server rooms as software that inspects data at rest.

Most cloud providers offer cloud-based DLP solutions to protect against data breaches and misuse of data. These solutions often integrate with software, infrastructure, and platform services and can include any of the systems mentioned previously. Cloud-based DLP is necessary for companies that have increased bring-your-own-device (BYOD) usage, and that store data and operate infrastructure within the cloud.

As with host intrusion detection systems (HIDS) and network intrusion detection systems (NIDS), DLP solutions must be accurate and updated to reduce the number of false positives and false negatives. Most systems alert the security administrator if there is a possibility of data leakage. However, it is up to the administrator to determine whether the threat is real.

Secure Sockets Layer (SSL)/Transport Layer Security (TLS) Inspection

Key Topic

Transport Layer Security Inspection (TLSI) is a security process that allows organizations to decrypt traffic, inspect the decrypted content for threats, and then reencrypt the traffic before it enters or leaves the network. TLSI is also known as "TLS break and inspect." TLS replaces the Secure Sockets Layer (SSL) protocol. In the past, TLSI was referred to as ***SSL Inspection (SSLI)***. Because newer implementations use TLS instead of SSL, this chapter uses the term *TLSI* to refer to this type of inspection. Using TLSI enhances visibility within some security products (such as IDS and DLP systems at the network edge) but also introduces new threats.

One of the main threats introduced when using TLSI is the potential abuse of the certificate authority (CA). Attackers can use this CA to issue unauthorized certificates trusted by the TLS clients. Abuse of a trusted CA could allow an attacker to sign malicious code to security controls and monitoring capabilities or to deploy malicious services that impersonate legitimate services to endpoints.

Some organizations can configure a policy to enforce traffic to be decrypted and inspected only "as authorized" to ensure that decrypted traffic is contained in an out-of-band and isolated segment of the network. This technique can be used to prevent unauthorized access to the decrypted traffic.

TIP As a mitigation, you should break and inspect TLS traffic only once within the organization's network. Redundant TLSI (where traffic is decrypted and inspected more than once) should not be performed. Inspecting multiple times can greatly complicate the ability to diagnose network issues with TLS traffic.

API Considerations

Application programming interfaces (APIs) are used in most modern applications. Organizations should evaluate different API considerations to make sure that APIs and underlying implementations are secure. If not configured or managed correctly, they can expose sensitive data to attackers. Furthermore, in the past there were several major incidents where the privacy of individuals was compromised. For example, several years ago Cambridge Analytica used Facebook's APIs to obtain information about millions of Facebook users. The misuse of similar APIs has sparked many privacy concerns from consumers and lawmakers. Nowadays, APIs are used in mobile apps, games, social networking platforms, dating apps, news websites, e-commerce sites, video and music streaming services, mobile payment systems, and many other implementations. When a third-party organization (such as an app developer or advertiser) obtains access to data through APIs, it may also gain access to very sensitive personal information. Some of these past events also sparked the creation of enhanced privacy laws in the United States, Europe, and worldwide.

Data Masking and Obfuscation

Key Topic

Data masking (otherwise known as data obfuscation) is the act of hiding sensitive information/data with specific characters or other data in order to protect it. For example, data masking (and obfuscation) has been used to protect personally identifiable information (PII) or commercially sensitive data from unauthorized users and attackers.

The following techniques have been used for data masking and obfuscation:

Key Topic

- **Substitution:** In this technique you substitute the original data with another authentic-looking value. For instance, you may replace a Social Security number or credit card number with a fake number. Figure 9-1 illustrates a basic example of substitution.

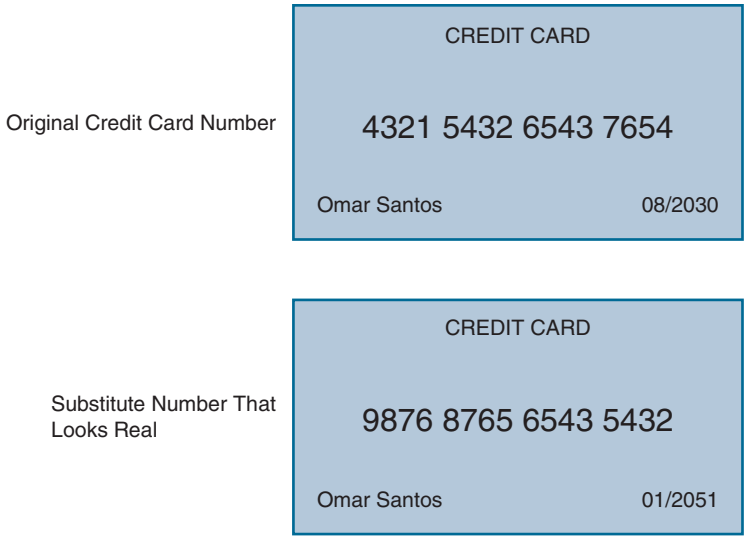


FIGURE 9-1 Data Masking Using Substitution

Figure 9-2 shows another example of masking by revealing only the last four digits of the credit card.

Key Topic

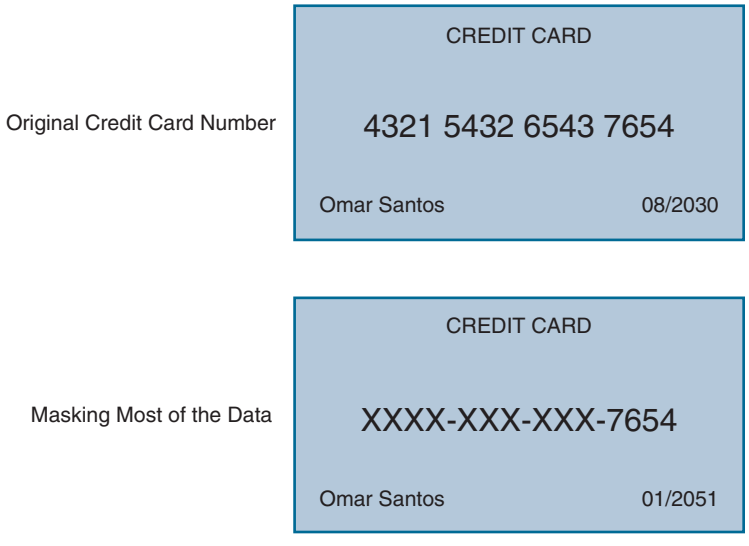


FIGURE 9-2 Additional Data Masking Example

- ***Tokenization***: This technique is used mostly when protecting data at rest. It is the process of randomly generating a token value for plaintext data and storing the mapping in a database. Tokenization is difficult to scale securely due to the performance and size of the underlying database. Tokenized data is difficult to exchange, since it requires direct access to the database (or token vault). Encryption provides a better way to scale and exchange sensitive data in a secure manner.

Encryption at Rest, in Transit/Motion, and in Processing

What do you need to encrypt? Without a doubt, you need to encrypt data, but more specifically three types of data: data in use, data at rest, and data in transit.

Figure 9-3 illustrates the concepts of encrypting data at rest, in transit/motion, and in use/processing.

Key Topic

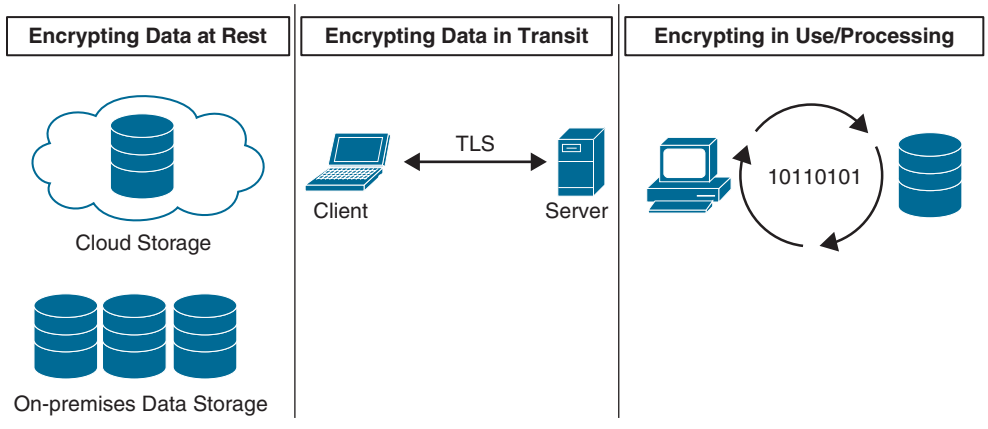


FIGURE 9-3 Encrypting Data at Rest, in Transit/Motion, and in Use/Processing

Key Topic

Data in use/processing can be described as actively used data undergoing constant change; for example, it could be stored in databases or spreadsheets. ***Data at rest*** is inactive data that is archived—backed up or stored in cloud storage services. ***Data in transit*** (also known as data in motion) is data that crosses the network or data that currently resides in computer memory.

Hashing

Key Topic

A ***hash*** is a summary of a file or message, often in numeric format. Hashes are used in digital signatures, in file and message authentication, and as a way to protect the integrity of sensitive data—for example, data entered into databases or perhaps

entire hard drives. A hash is generated through the use of a hash function to verify the integrity of the file or message, most commonly after transit over a network. A *hash function* is a mathematical procedure that converts a variable-sized amount of data into a smaller block of data. The hash function is designed to take an arbitrary data block from the file or message, use that as an input, and from that block produce a fixed-length hash value. Basically, the hash is created at the source and is recalculated and compared with the original hash at the destination.

Because the hash is a condensed version of the file/message, or a portion of it, it is also known as a message digest. It provides integrity to data so that a user knows that the message is intact, hasn't been modified during transit, and comes from the source the user expects. A hash can fall into the category of a *one-way function*. This means it is easy to compute when generated but difficult (or impossible) to compute in reverse. In the case of a hash, a condensed version of the message, initial computation is relatively easy (compared to other algorithms), but the original message should not be re-created from the hash. Contrast this concept to encryption methods that indeed can be reversed. A hash can be created without the use of an algorithm, but generally, the ones used in the field require some kind of cryptographic algorithm.

NOTE In Chapter 16, "Summarizing the Basics of Cryptographic Concepts," you will learn the details about different hashing algorithms such as MD5 and different versions of SHA. You will also learn which flavors of those algorithms are recommended and which should be avoided (such as MD5).

Rights Management

Key Topic

Digital rights management (DRM) is the name given to a set of access control technologies that are used to control the use of proprietary hardware, software, and copyrighted works. DRM solutions are used to restrict the use, modification, and distribution of copyrighted works and the underlying systems used to enforce such policies.

Traditional DRM policies include restrictive licensing agreements and solutions that encrypt expressive material or embed a digital tag designed to control access and reproduction of information. Software companies and content publishers typically enforce their own access policies on content (that is, restrictions on copying or viewing such content or software). Hardware manufacturers have also expanded the usage of DRM to more traditional hardware products.

Digital rights management policies and processes are not universally accepted. Some argue that there is no substantial evidence that DRM helps to completely prevent copyright infringement. The following section covers some of the geographical considerations for data protection and DRM.

Geographical Considerations

Many ***geographical considerations*** (depending on where you reside) affect the laws and regulations that have been created to address data privacy and DRM:

- **Data privacy:** Earlier in this chapter, you learned about GDPR (a regulation in the European Union [EU] and the European Economic Area [EEA] focused on data protection and privacy). You also learned about the California Consumer Privacy Act (CCPA). Many other laws and regulations in other countries give consumers the right to know what personal information is being collected by companies, government, and any other organizations. Additional examples include the Australia Privacy Principles (APP), Canada's Personal Information Protection and Electronic Data Act (PIPEDA), and Japan's Personal Information Protection Act (JPIPA).
- **DRM:** Many laws around the world have been created to criminalize the circumvention of DRM, communication about such circumvention, and the creation and distribution of tools used. Examples of these DRM-related laws are the United States Digital Millennium Copyright Act (US-DMCA) and the European Union's Information Society Directive.

TIP Different countries might have laws that dictate how personal data is stored and transferred between different geographical locations and countries. In addition, some regulations (such as GDPR) define policies for "third countries" and what personal data can or cannot be transferred between countries. Additional information about GDPR "third countries" definitions and information can be obtained from https://gdpr-info.eu/issues/third-countries/.

Data Breach Response and Recovery Controls

You must have the proper ***response and recovery controls*** in place in the unfortunate event of a data breach. This data breach response and recovery control plan must include the assembly of a team of experts within your organization, as well as legal counsel. You also have to identify a data forensics team (in most cases hired from third-party providers). The forensics team must help determine the source and scope of the breach, as well as collect and analyze evidence to outline remediation steps.

In the case of the data breach, you should notify law enforcement (when applicable) and also notify the affected businesses and individuals based on local and federal laws.

TIP The United States Federal Trade Commission (FTC) provides detailed guidance and recommendations for post-breach response and recovery at www.ftc.gov/tips-advice/business-center/guidance/data-breach-response-guide-business.

Site Resiliency

Within the confidentiality, integrity, availability (CIA) triad, redundant sites fall into the category of *availability*. In the case of a disaster, a redundant site can act as a safe haven for your data and users. Redundant sites are sort of a gray area between redundancy and a disaster recovery method. If you have one and need to use it, a "disaster" has probably occurred. But the better the redundant site, the less time the organization loses, and the less it seems like a disaster and more like a failure that you have prepared for. Of course, this outcome all depends on the type of redundant site your organization decides on.

Regarding the types of redundant sites, I like to refer to the story of Goldilocks and the three bears' three bowls of porridge. One was too hot, one too cold, and one just right. Most organizations opt for the warm redundant site as opposed to the hot or cold. Let's look at these three now.

Key Topic

- ***Hot site*:** This site is a near duplicate of the original site of the organization that can be up and running within minutes (maybe longer). Computers and phones are installed and ready to go, a simulated version of the server room stands ready, and the vast majority of the data is replicated to the site on a regular basis in the event that the original site is not accessible to users for whatever reason. Hot sites are used by companies that would face financial ruin in the case that a disaster makes their main site inaccessible for a few days or even a few hours. This is the only type of redundant site that can facilitate a full recovery.
- ***Warm site*:** This site has computers, phones, and servers, but they might require some configuration before users can start working on them. The warm site will have backups of data that might need to be restored; they will probably be several days old. This type of site is chosen the most often by organizations because it has a good amount of configuration yet remains less expensive than a hot site.

- ***Cold site*:** This site has tables, chairs, bathrooms, and possibly some technical setup—for example, basic phone, data, and electric lines. Otherwise, a lot of configuration of computers and data restoration is necessary before the site can be properly utilized. This type of site is used only if a company can handle the stress of being nonproductive for a week or more.

Although they are redundant, these types of sites are generally known as backup sites because if they are required, a disaster has probably occurred. A good network security administrator tries to plan for, and rely on, redundancy and fault tolerance as much as possible before having to resort to disaster recovery methods.

Deception and Disruption

Honeypots and honeynets attract and trap potential attackers to counteract any attempts at unauthorized access of the network. These solutions isolate the potential attacker in a monitored area and contain dummy resources that look to be of value to the perpetrator. While an attacker is trapped in one of these, the attacker's methods can be studied and analyzed, and the results of those analyses can be applied to the general security of the functional network.

Key Topic

A ***honeypot*** is generally a single computer but could also be a file, group of files, or an area of unused IP address space, whereas a ***honeynet*** is one or more computers, servers, or an area of a network; a honeynet is used when a single honeypot is not sufficient. Either way, the individual computer, or group of servers, will *usually* not house any important company information. Various analysis tools are implemented to study the attacker; these tools, along with a centralized group of honeypots (or a honeynet), are known collectively as a honeyfarm.

One example of a honeypot in action is the spam honeypot. Spam email is one of the worst banes known to network administrators; a spam honeypot can lure spammers in, enabling network administrators to study the spammers' techniques and habits, thus allowing the network admins to better protect their actual email servers, SMTP relays, SMTP proxies, and so on, over the long term. This solution might ultimately keep the spammers away from the real email addresses because the spammers are occupied elsewhere. Some of the information gained by studying spammers is shared with other network admins or organizations' websites dedicated to reducing spam. A spam honeypot could be as simple as a single email address or as complex as an entire email domain with multiple SMTP servers.

Of course, as with any technology that studies attackers, honeypots also bear risks to the legitimate network. The honeypot or honeynet should be carefully firewalled off from the legitimate network to ensure that the attacker can't break through.

Key Topic

Honeyfiles can also be used as bait files intended to lure adversaries to access and then send alarms to security analysts for detection. They can also be used to potentially learn the tactics and techniques used by attackers. For instance, you can create a honeyfile named credentials.txt to lure attackers to access it. Honeyfiles can be used to learn the attackers' tactics, techniques, and behavior without adversely affecting normal operations.

Fake Telemetry

Key Topic

Some organizations have used additional deception techniques, such as ***fake telemetry***, as decoys and breadcrumbs in order to lure and trick attackers. Similarly, attackers have compromised systems that will also generate fake telemetry and reporting data to fool security monitoring systems, analysts in a security operations center (SOC), and evade other security controls that may be in place.

TIP The MITRE ATT&CK framework includes examples of these deception and evasion techniques. Examples include the T0856 technique described at https://collaborate.mitre.org/attackics/index.php/Technique/T0856.

DNS Sinkhole

Key Topic

Another deception and disruption technique is the use of ***DNS sinkholes***, or "blackhole DNS servers." In a DNS sinkhole, you configure one or more DNS servers to provide false results to attackers and redirect them to areas in the network where you can observe their tactics and techniques. DNS sinkholes have been used to contain different types of malware such as the infamous WannaCry ransomware and to disrupt certain malicious DNS operations in denial-of-service (DoS) and other attacks. For example, these DNS sinkholes have been used to interrupt DNS resolution to malicious command and control (C2) servers and botnet coordination.

TIP Adversaries have also used similar techniques to perform DNS poisoning attacks to redirect systems and users to malicious destinations.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 9-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 9-2 Key Topics for Chapter 9

Key Topic Element	Description	Page Number
List	Listing the key elements used in the process of configuration management	213
Paragraph	Defining data sovereignty	214
Paragraph	Defining data loss prevention (DLP) systems	214
Paragraph	Understanding Transport Layer Security Inspection	215
Paragraph	Defining data masking and obfuscation	216
List	Listing the different techniques used for data masking	216
Figure 9-2	Additional data masking example	217
Figure 9-3	Encrypting data at rest, in transit/motion, and in use/processing	218
Paragraph	Understanding encryption of data in processing, data at rest, and data in transit/motion	218
Paragraph	Defining hashing and hashing functions	218
Paragraph	Defining digital rights management	219
List	Defining hot site, warm site, and cold site in the context of site resiliency	221
Paragraph	Defining honeypots and honeynets	222
Paragraph	Describing honeyfiles	223
Paragraph	Understanding the use of fake telemetry	223
Paragraph	Understanding the use of DNS sinkholes	223

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

configuration management, diagrams, baseline configuration, standard naming conventions, Internet Protocol (IP) schema, data sovereignty, data protection, data loss prevention (DLP), Transport Layer Security Inspection (TLSI), SSL Inspection (SSLI), data masking, tokenization, data in use/processing, data at rest, data in transit, hash, digital rights management (DRM), geographical considerations, response and recovery controls, hot site, warm site, cold site, honeypot, honeynet, honeyfiles, fake telemetry, DNS sinkholes

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. In the context of site resiliency, a _____ will have backups of data that might need to be restored; they will probably be several days old. This type of site is chosen most often by organizations because it has a good amount of configuration yet remains less expensive than a hot site.
2. What can be used as bait files intended to lure adversaries to access and then send alarms to security analysts for detection?
3. What is the name given to a set of access control technologies that are used to control the use of proprietary hardware, software, and copyrighted works?
4. What term is used when data is actively used and undergoing constant change? For instance, data could be stored in databases or spreadsheets and be processed by running applications.
5. What is the process of generating a random value for plaintext data and storing the mapping in a database?



This chapter covers the following topics related to Objective 2.2 (Summarize virtualization and cloud computing concepts) of the CompTIA Security+ SY0-601 certification exam:

- Cloud Models
  - Infrastructure as a service (IaaS)
  - Platform as a service (PaaS)
  - Software as a service (SaaS)
  - Anything as a service (XaaS)
  - Public
  - Community
  - Private
  - Hybrid
- Cloud service providers
- Managed service provider (MSP)/ managed security service provider (MSSP)
- On-premises vs. off-premises
- Fog computing
- Edge computing
- Thin client
- Containers
- Microservices/API
- Infrastructure as code
  - Software-defined networking (SDN)
  - Software-defined visibility (SDV)
- Serverless architecture
- Services integration
- Resource policies
- Transit gateway
- Virtualization
  - Virtual machine (VM) sprawl avoidance
  - VM escape protection

Summarizing Virtualization and Cloud Computing Concepts

These days, more and more organizations are transferring some or all of their server and network resources to the cloud. This effort creates many potential hazards and vulnerabilities that you, as the security administrator, and the cloud provider must address. Top among these concerns are the servers, where all data is stored and accessed. Servers of all types should be hardened and protected from a variety of attacks to keep the integrity of data from being compromised. However, data must also be available. Therefore, you must strike a balance of security and availability. This chapter provides details about several virtualization and cloud computing concepts.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 10-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 10-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Cloud Models	1–3
Cloud Service Providers	4–5
Cloud Architecture Components	6–8
Virtual Machine (VM) Sprawl Avoidance and VM Escape Protection	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following cloud service models will you use if you want to host applications on virtual machines, deploy load balancers, and use storage buckets?
   a. IaaS
   b. PaaS
   c. SaaS
   d. None of these answers are correct.

2. Which of the following cloud deployments is a mix of public and private cloud solutions where multiple organizations can share the public cloud portion?
   a. Community cloud
   b. PaaS
   c. SaaS
   d. MSSP cloud

3. Google Drive, Office 365, and Dropbox are examples of which of the following types of cloud service?
   a. SaaS
   b. IaaS
   c. PaaS
   d. All of these answers are correct.

4. What type of company or organization provides services to manage your security devices and can also help monitor and respond to security incidents?
   a. SaaS provider
   b. PaaS provider
   c. MSSP
   d. Serverless provider

5. Which of the following organizations delivers network, application, system, and management services using a pay-as-you-go model?
   a. SaaS provider
   b. PaaS provider
   c. XaaS provider
   d. Managed service provider (MSP)

6. Which term is used to describe an ecosystem of resources and applications in new network services (including 5G and IoT)?
   a. SaaS
   b. VPC
   c. Edge computing
   d. None of these answers are correct.

7. Which of the following are computer systems that run from resources stored on a central server or from the cloud instead of a local (on-premises) system?
   a. Thin clients
   b. Fog edge devices
   c. VPCs
   d. Containers

8. Which of the following are technologies and solutions to manage, deploy, and orchestrate containers?
   a. Docker Swarm
   b. Apache Mesos
   c. Kubernetes
   d. All of these answers are correct.

9. What condition could occur when an organization can no longer effectively control and manage all the VMs on a network or in the cloud?
   a. VM sprawl
   b. VM escape
   c. Hypervisor escape
   d. Hypervisor sprawl

10. Which condition occurs when an attacker or malware compromises one VM and then attacks the hypervisor?
    a. VM escape
    b. Hypervisor escape
    c. VM sprawl
    d. Hypervisor sprawl

Foundation Topics

Cloud Models

Cloud computing can be defined as a way of offering on-demand services that extend the capabilities of a person's computer or an organization's network. These might be free services, such as personal browser-based email from various providers, or they could be offered on a pay-per-use basis, such as services that offer data access, data storage, infrastructure, and online gaming. A network connection of some sort is required to make the connection to the cloud and gain access to these services in real time.

Some of the benefits to an organization using cloud-based services include lowered cost, less administration and maintenance, more reliability, increased scalability, and possible increased performance. A basic example of a cloud-based service would be browser-based email. A small business with few employees definitely needs email, but it can't afford the costs of an email server and perhaps does not want to have its own hosted domain and the costs and work that go along with that. By connecting to a free web browser-based service, the small business can obtain near unlimited email, contacts, and calendar solutions. However, there is no administrative control, and there are some security concerns, which are discussed in a little bit.

Cloud computing services are generally broken down into several categories of services:

Key Topic

- ***Software as a service (SaaS)*:** The most commonly used and recognized of the cloud service categories, SaaS offers users access to applications that are provided by a third party over the Internet. The applications need not be installed on the local computer. In many cases these applications are run within a web browser; in other cases the user connects with screen-sharing programs or remote desktop programs. A common example is webmail.

NOTE Often compared to SaaS is the application service provider (ASP) model. SaaS typically offers a generalized service to many users. However, an ASP typically delivers a service (perhaps a single application) to a small number of users.

- ***Infrastructure as a service (IaaS)*:** This service offers computer networking, storage, load balancing, routing, and VM hosting. More and more organizations are seeing the benefits of offloading some of their networking infrastructure to the cloud.

- ***Platform as a service (PaaS)*:** This service provides various software solutions to organizations, especially the ability to develop applications in a virtual environment without the cost or administration of a physical platform. PaaS is used for easy-to-configure operating systems and on-demand computing. Often, this utilizes IaaS as well for an underlying infrastructure to the platform. Cloud-based virtual desktop environments (VDEs) and virtual desktop infrastructures (VDIs) are often considered to be part of this service but can be part of IaaS as well.
- ***Anything as a service (XaaS)*:** There are many types of cloud-based services. If they don't fall into the previous categories, they often fall under the category "anything as a service," or XaaS. For instance, a service in which a large service provider integrates its security services into the company's/customer's existing infrastructure is often referred to as security as a service (SECaaS). The concept is that the service provider can provide the security more efficiently and more cost effectively than a company can, especially if it has a limited IT staff or budget. The Cloud Security Alliance (CSA) defines various categories to help businesses implement and understand SECaaS, including encryption, data loss prevention (DLP), continuous monitoring, business continuity and disaster recovery (BCDR), vulnerability scanning, and much more. Periodically, new services will arrive, such as monitoring as a service (MaaS)—a framework that facilitates the deployment of monitoring within the cloud in a continuous fashion.

Key Topic

TIP NIST's special publication (SP) 800-145 provides a great overview of the definitions of cloud computing concepts. You can access this document at https://csrc.nist.gov/publications/detail/sp/800-145/final.

Public, Private, Hybrid, and Community Clouds

Organizations use different types of clouds: public, private, hybrid, and community.

Key Topic

- ***Public cloud*:** In this type of cloud, a service provider offers applications and storage space to the general public over the Internet. Examples include free, web-based email services and pay-as-you-go business-class services. The main benefits of this type of service include low (or zero) cost and scalability. Providers of public cloud space include Google, Rackspace, and Amazon.
- ***Private cloud*:** This type of cloud is designed with a particular organization in mind. In it, you, as security administrator, have more control over the data and infrastructure. A limited number of people have access to the cloud, and they are usually located behind a firewall of some sort in order to gain

access to the private cloud. Resources might be provided by a third-party or could come from your server room or data center. In other words, the private cloud could be deployed using resources in an on-premises data center or dedicated resources by a third-party cloud provider. An example is Amazon AWS Virtual Private Cloud (https://aws.amazon.com/vpc) and Microsoft Azure private cloud offering (https://azure.microsoft.com/en-us/overview/what-is-a-private-cloud).

- ***Hybrid cloud*:** This type is a mixture of public and private clouds. Dedicated servers located within the organization and cloud servers from a third party are used together to form the collective network. In these hybrid scenarios, confidential data is usually kept in-house.
- ***Community cloud*:** This is another mix of public and private cloud deployments where multiple organizations can share the public portion. Community clouds appeal to organizations that usually have a common form of computing and storing of data.

The type of cloud an organization uses is dictated by its budget, the level of security it requires, and the amount of manpower (or lack thereof) it has to administer its resources. Although a private cloud can be very appealing, it is often beyond the ability of an organization, forcing that organization to seek a public or community-based cloud. However, it doesn't matter what type of cloud is used. Resources still have to be secured by someone, and you'll have a hand in that security one way or the other.

Cloud Service Providers

A ***cloud service provider (CSP)***, such as Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform (GCP), might offer one or more of the aforementioned services. CSPs have no choice but to take their security and compliance responsibilities very seriously. For instance, Amazon created a Shared Responsibility Model to define in detail the responsibilities of AWS customers as well as those of Amazon. You can access the Amazon Shared Responsibility Model at https://aws.amazon.com/compliance/shared-responsibility-model.

NOTE The shared responsibility depends on the type of cloud model, whether SaaS, PaaS, or IaaS.

Similar to the CSP is the ***managed service provider (MSP)***, which can deliver network, application, system, and management services using a pay-as-you-go model. An MSP is an organization that can manage your network infrastructure, servers,

and in some cases your security devices. Those companies that provide services to manage your security devices and can also help monitor and respond to security incidents are called ***managed security service providers (MSSPs)***.

TIP Another type of security service provided by vendors is *managed detection and response (MDR)*. MDR evolved from MSSPs to help organizations that lack the resources to establish robust security operations programs to detect and respond to threats. MDR is different from traditional MSSP services because it is more focused on threat detection rather than compliance and traditional network security device configuration services. Most of the services offered by MDR providers are delivered using their own set of tools and technologies but are deployed on the customer's premises. The techniques MDR providers use may vary. A few MDR providers rely solely on security logs, whereas others use network security monitoring or endpoint activity to secure the network. MDR solutions rely heavily on security event management; Security Orchestration, Automation, and Response (SOAR); and advanced analytics.

Cloud Architecture Components

Cloud computing allows you to store and access your data, create applications, and interact with thousands of other services over the Internet (***off-premises***) instead of your local network and computers (***on-premises***). With a good Internet connection, cloud computing can be executed anytime, anywhere. There are several cloud-related architecture components, described next:

- The concept of fog and edge computing
- Thin clients
- Containers
- Microservices and related APIs

Fog and Edge Computing

Key Topic

The two architectural concepts called fog computing and edge computing are often used interchangeably. However, there are a few differences. The term ***edge computing*** describes an ecosystem of resources and applications in new network services (including 5G and IoT). One of the main benefits is to provide greater network speeds, low latency, and computational power near the user.

Figure 10-1 illustrates the concept of edge computing.

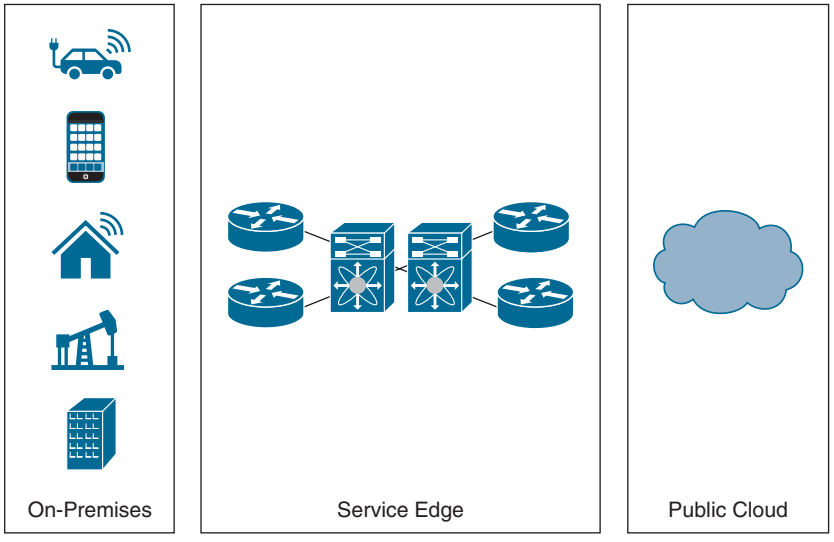


FIGURE 10-1 Edge Computing

The "edge" in edge computing is the location nearer the consumer (otherwise known as the subscriber) and where data is processed or stored without being back-hauled to a central location. This service edge is optimized to provide a low-cost solution, increase performance, and provide a better user experience.

Key Topic

The term ***fog computing*** was originally created by Cisco in the early 2010s. It describes the decentralization of computing infrastructure by "bringing the cloud to the ground" (thus the term *fog*). This architecture enables components of the edge computing concept to easily push compute power away from the public cloud to improve scalability and performance. Fog computing accomplishes this by putting data, compute, storage, and applications near the end user or IoT device.

Other similar and related concepts have been adopted throughout the years. For example, consider the multi-access edge computing (MEC) concept created by the European Telecommunications Standards Institute (ETSI), designed to benefit application developers and content providers. Another term used by some organizations is *cloudlets*, which are micro data centers that can be deployed locally near the data source (in some cases temporarily for emergency response units and other use cases).

Thin Clients

Key Topic

Thin clients are computer systems that run from resources stored on a central server or from the cloud instead of a local (on-premises) system. When you use a thin

client, you connect remotely to a server-based computing environment where the applications, sensitive data, and memory are stored.

Thin clients can be used in several ways:

- **Shared terminal services:** Users at thin client stations share a server-based operating system and applications.
- **Desktop virtualization:** A full desktop machine (including the operating system and applications) lives in a virtual machine in an on-premises server or in the cloud (off-premises).
- **Browser-based thin clients:** All applications are accessed within a web browser.

Figure 10-2 shows a user accessing a virtual desktop running Microsoft Windows and custom applications running in the cloud.

Key Topic

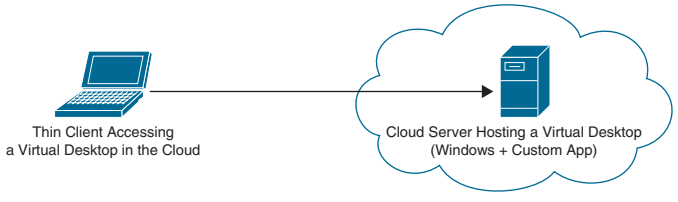


FIGURE 10-2 Thin Client Accessing a Virtualized Desktop

Containers

Key Topic

Before you can even think of building a distributed system, you must first understand how the container images that contain your applications make up all the underlying pieces of such a distributed system. Applications are normally composed of a language runtime, libraries, and source code. For instance, your application may use third-party or open-source shared libraries such as the Linux Kernel (www.kernel.org), nginx (https://nginx.org), or OpenSSL (www.openssl.org). These shared libraries are typically shipped as shared components in the operating system that you installed on a system. The dependency on these libraries introduces difficulties when an application developed on your desktop, laptop, or any other development machine (dev system) has a dependency on a shared library that isn't available when the program is deployed out to the production system. Even when the dev and production systems share the exact same version of the operating system, bugs can occur when programmers forget to include dependent asset files inside a package that they deploy to production.

The good news is that you can package applications in a way that makes it easy to share them with others. In this case, ***containers*** become very useful. Docker, one of the most popular container runtime engines, makes it easy to package an executable and push it to a remote registry where it can later be pulled by others.

Container registries are available in all of the major public cloud providers (for example, AWS, Google Cloud Platform, and Microsoft Azure) as well as services to build images. You can also run your own registry using open-source or commercial systems. These registries make it easy for developers to manage and deploy private images, while image-builder services provide easy integration with continuous delivery systems.

Container images bundle a program and its dependencies into a single artifact under a root file system. Containers are made up of a series of file system layers. Each layer adds, removes, or modifies files from the preceding layer in the file system. The overlay system is used both when packaging up the image and when the image is actually being used. During runtime, there are various different concrete implementations of such file systems, including aufs, overlay, and overlay2.

Let's look at an example of how container images work. Figure 10-3 shows three container images: A, B, and C. Container Image B is "forked" from Container Image A. Then, in Container Image B, Python version 3 is added. Furthermore, Container Image C is built on Container Image B, and the programmer adds OpenSSL and nginx to develop a web server and enable TLS.

Key Topic

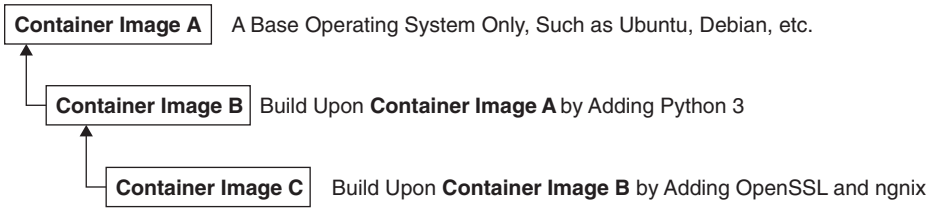


FIGURE 10-3 How Container Images Work

Abstractly, each container image layer builds on the previous one. Each parent reference is a pointer. The example in Figure 10-3 includes a simple set of containers; in many environments, you will encounter a much larger directed acyclic graph.

Even though the Security+ exam does not cover Docker in detail, it is still good to see a few examples of Docker containers, images, and related commands. Figure 10-4 shows the output of the **docker images** command.

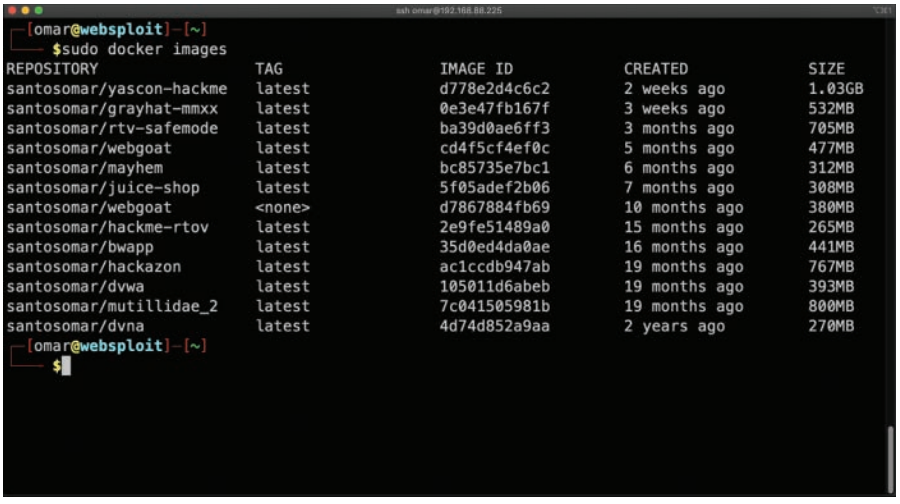


FIGURE 10-4 Displaying the Docker Images

The Docker images shown in Figure 10-4 are intentionally vulnerable applications that you can also use to practice your skills. These Docker images and containers are included in a VM called WebSploit (websploit.org) by Omar Santos. The VM is built on top of Kali Linux and includes several additional tools, along with the aforementioned Docker containers. This can be a good tool, not only to get familiarized with Docker, but also to learn and practice offensive and defensive security skills.

Figure 10-5 shows the output of the **docker ps** command used to see all the running Docker containers in a system.

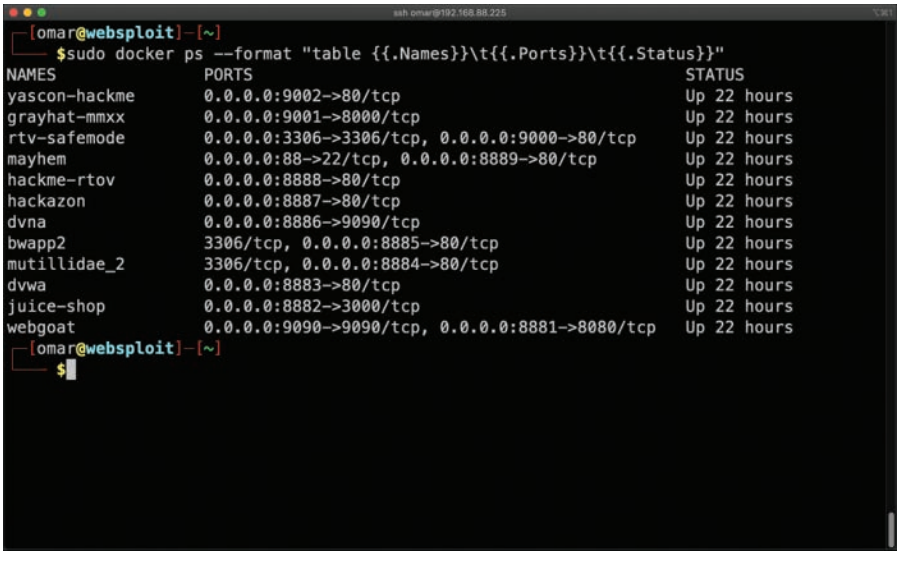


FIGURE 10-5 Output of the **docker ps** Command

You can use a public, cloud provider, or private Docker image repository. Docker's public image repository is called Docker Hub (https://hub.docker.com). You can find images by going to the Docker Hub website or by using the **docker search** command, as demonstrated in Figure 10-6.

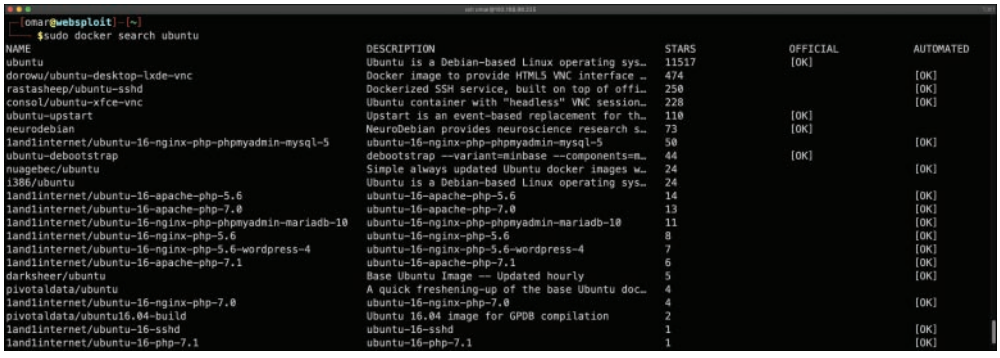


FIGURE 10-6 Output of the **docker search** Command

In Figure 10-6, the user searches for a container image that matches the **ubuntu** keyword.

TIP You can practice and deploy your first container by using Katacoda (an interactive system that allows you to learn many different technologies, including Docker, Kubernetes, Git, TensorFlow, and many others). You can access Katacoda at www.katacoda.com. There are numerous interactive scenarios provided by Katacoda. For instance, you can use the "Deploying Your First Docker Container" scenario to learn (hands-on) Docker: www.katacoda.com/courses/docker/deploying-first-container.

In larger environments, you will not deploy and orchestrate Docker containers in a manual way. You will want to automate as much as possible. This is where Kubernetes comes into play. Kubernetes (often referred to as k8s) automates the distribution, scheduling, and orchestration of application containers across a cluster.

The following are the Kubernetes components:

- **Control Plane:** This component coordinates all the activities in your cluster (scheduling, scaling, and deploying applications).
- **Node:** This VM or physical server acts as a worker machine in a Kubernetes cluster.

- **Pod:** This group of one or more containers provides shared storage and networking, including a specification for how to run the containers. Each pod has an IP address, and it is expected to be able to reach all other pods within the environment.

There have been multiple technologies and solutions to manage, deploy, and orchestrate containers in the industry. The following are the most popular:

- **Kubernetes:** One of the most popular container orchestration and management frameworks, originally developed by Google, Kubernetes is a platform for creating, deploying, and managing distributed applications. You can download Kubernetes and access its documentation at https://kubernetes.io.
- **Nomad:** This container management and orchestration platform is by HashCorp. You can download and obtain detailed information about it at www.nomadproject.io.
- **Apache Mesos:** This distributed Linux kernel provides native support for launching containers with Docker and AppC images. You can download Apache Mesos and access its documentation at https://mesos.apache.org.
- **Docker Swarm:** This container cluster management and orchestration system is integrated with the Docker Engine. You can access the Docker Swarm documentation at https://docs.docker.com/engine/swarm.

Microservices and APIs

Key Topic

The term ***microservices*** describes how applications can be deployed as a collection of services that are highly maintainable and testable and independently deployable. Microservices are organized around business capabilities and enable the rapid, frequent, and reliable delivery of large, complex applications. Most microservices developed and deployed by organizations are based on containers.

Application programming interfaces (APIs), such as RESTful APIs, are the frameworks through which developers can interact with an application. Microservices use APIs to communicate between themselves. These APIs can be used for inter-microservice communication, and they can also be used to expose data and application functionality to third-party systems.

Microservices and APIs need to be secured. Traditional segmentation strategies do not work well in these virtual and containerized environments. The ability to enforce network segmentation in container and VM environments is what people call *micro-segmentation*. Micro-segmentation is at the VM level or between containers regardless of a VLAN or a subnet. Micro-segmentation solutions need to be

"application aware." This means that the segmentation process starts and ends with the application itself.

Most micro-segmentation environments apply a zero-trust model. This model dictates that users cannot talk to applications and that applications cannot talk to other applications unless a defined set of policies permits them to do so.

Infrastructure as Code

Modern applications and deployments (on-premises and in the cloud) require scalability and integration with numerous systems, APIs, and network components. This complexity introduces risks, including network configuration errors that can cause significant downtime and network security challenges. Consequently, networking functions such as routing, optimization, and security have also changed. The next generation of hardware and software components in enterprise networks must support both the rapid introduction and rapid evolution of new technologies and solutions. Network infrastructure solutions must keep pace with the business environment and support modern capabilities that help drive simplification within the network.

These elements have fueled the creation of ***software-defined networking (SDN)***. SDN was originally created to decouple control from the forwarding functions in networking equipment. This is done to use software to centrally manage and "program" the hardware and virtual networking appliances to perform forwarding. Organizations are adopting a framework often referred to as ***infrastructure as code***, which is the process of managing and provisioning computer data centers through machine-readable definition files rather than physical hardware configuration or interactive configuration tools.

In traditional networking, three different "planes," or elements, allow network devices to operate: the management, control, and data planes. The control plane has always been separated from the data plane. There was no central brain (or controller) that controlled the configuration and forwarding. Routers, switches, and firewalls were managed by the command-line interface (CLI), graphical user interfaces (GUIs), and custom Tcl scripts.

NOTE Tcl is a high-level, general-purpose, interpreted, dynamic programming language. It was designed with the goal of being very simple but powerful. Tcl casts everything into the mold of a command, even programming constructs like variable assignment and procedure definition.

For instance, the firewalls were managed by standalone web portals, while the routers were managed by the CLI. SDN introduced the notion of a *centralized controller*. The SDN controller has a global view of the network, and it uses a common management protocol to configure the network infrastructure devices.

In the topology shown in Figure 10-7, an SDN controller is used to manage a network infrastructure in a data center.

Key Topic

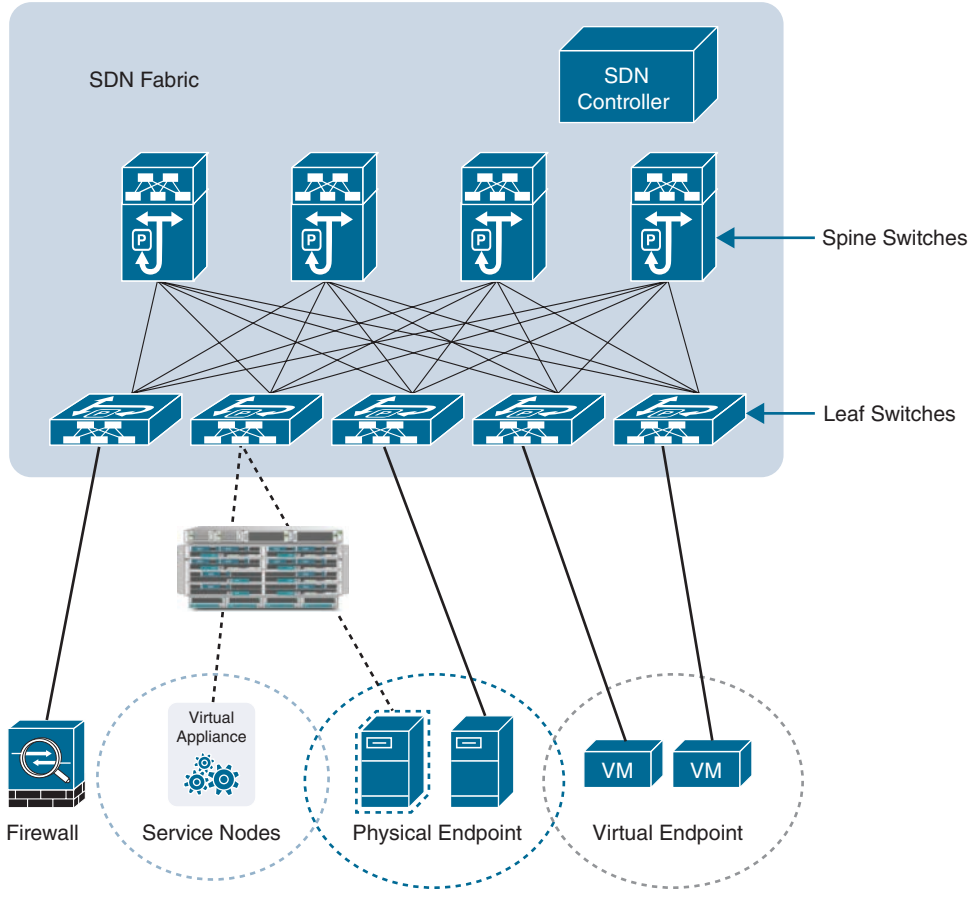


FIGURE 10-7 A High-Level SDN Implementation

The SDN controller can also calculate reachability information from many systems in the network, and it pushes a set of flows inside the switches. The hardware uses the flows to do the forwarding. Here, you can see a clear transition

from a distributed "semi-intelligent brain" approach to a "central and intelligent brain" approach.

TIP An example of an open-source implementation of SDN controllers is the Open vSwitch (OVS) project using the OVS Database (OVSDB) management protocol and the OpenFlow protocol. Another example is the Cisco Application Policy Infrastructure Controller (Cisco APIC). Cisco APIC is the main architectural component and the brain of the Cisco Application Centric Infrastructure (ACI) solution. OpenDaylight (ODL) is another popular open-source project that is focused on the enhancement of SDN controllers to provide network services across multiple vendors. OpenDaylight participants also interact with the OpenStack Neutron project and attempt to solve the existing inefficiencies. OpenDaylight interacts with Neutron via a northbound interface and manages multiple interfaces southbound, including the OVSDB and OpenFlow.

SDN changed a few things in the management, control, and data planes. However, the big change was in the control and data planes in software-based switches and routers (including virtual switches inside hypervisors). For instance, the Open vSwitch project started some of these changes across the industry.

SDN provides numerous benefits in the area of the management plane. These benefits are in both physical switches and virtual switches. SDN is now widely adopted in data centers.

Another concept introduced a few years ago is ***software-defined visibility (SDV)***. SDV is a concept similar to SDN but is designed and optimized to provide intelligent visibility of what's happening across an organization network or in the cloud. Typically, RESTful APIs are used in an SDV environment to orchestrate and automate security operations and monitor tasks across the organization.

Serverless Architecture

Another popular architecture is the ***serverless architecture***. Be aware that *serverless* does not mean that you do not need a "server" somewhere. Instead, it means that you will be using cloud platforms to host and/or to develop your code. For example, you might have a serverless app that is distributed in a cloud provider like AWS, Azure, or Google Cloud Platform.

Serverless is a cloud computing execution model where the cloud provider (AWS, Azure, Google Cloud, and so on) dynamically manages the allocation and provisioning of servers. Serverless applications run in stateless containers that are ephemeral and event-triggered (fully managed by the cloud provider). AWS Lambda is one of the most popular serverless architectures in the industry. Figure 10-8 shows a "function" or application in AWS Lambda.

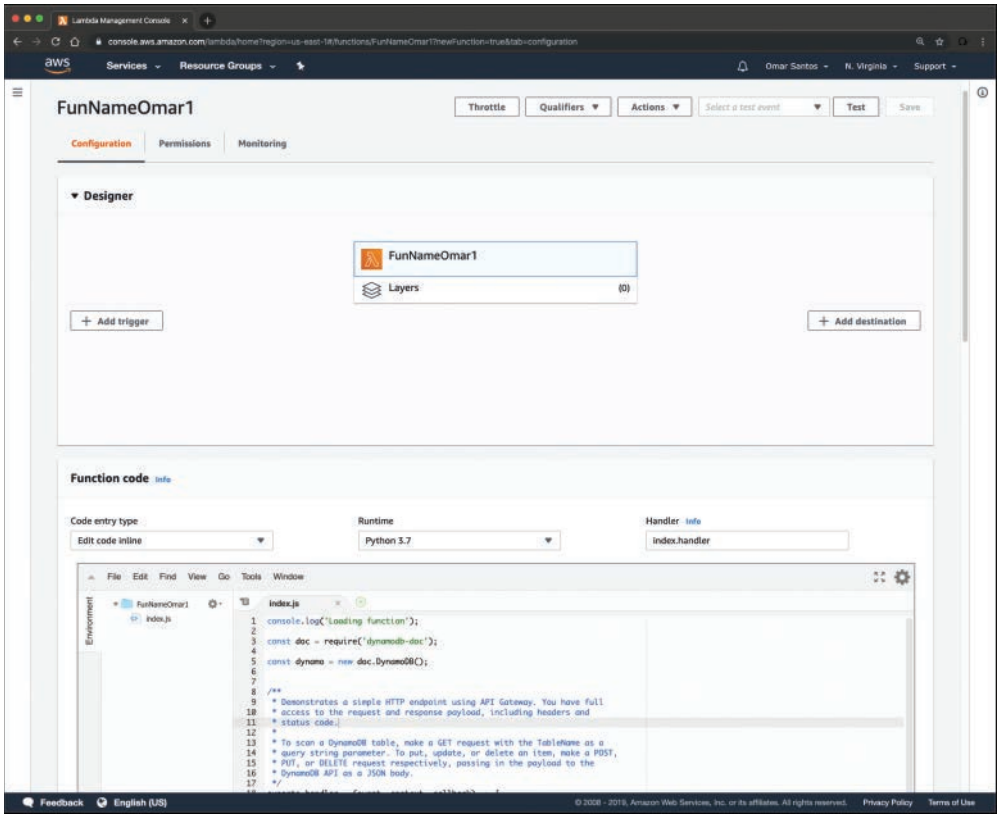


FIGURE 10-8 AWS Lambda

In AWS Lambda, you run code without provisioning or managing servers, and you pay only for the compute time you consume. When you upload your code, Lambda takes care of everything required to run and scale your application (offering high availability and redundancy).

Figure 10-9 summarizes the evolution of computing from physical servers to virtual machines, containers, and serverless solutions. Virtual machines and containers could be deployed on-premises or in the cloud.

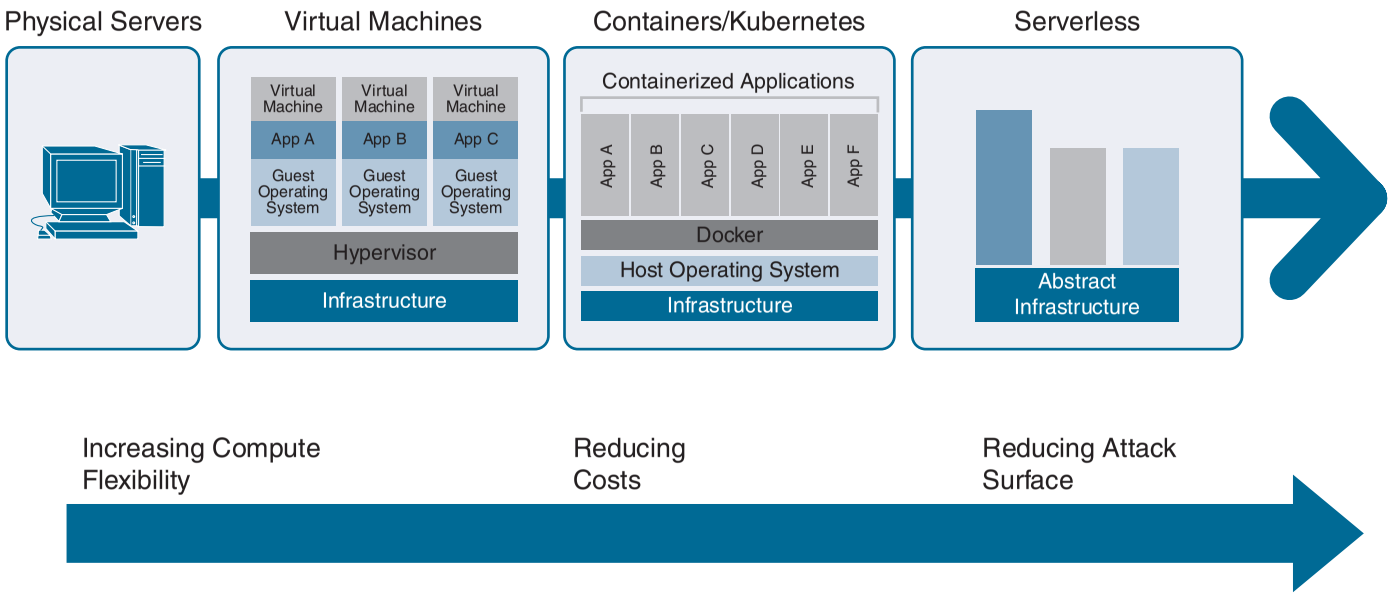


FIGURE 10-9 The Evolution of Computing from Physical Servers to Serverless

Services Integration

Key Topic

Cloud ***services integration*** is a series of tools and technologies used to connect different systems, applications, code repositories, and physical or virtual network infrastructure to allow real-time exchange of data and processes. Subsequently, integrated cloud services can be accessed by multiple systems, the Internet, or in the case of a private cloud, over a private network.

One of the main goals of cloud services integration is to break down data silos, increase visibility, and improve business processes. These services integrations are crucial for most organizations needing to share data among cloud-based applications. Cloud services integration has grown in popularity over the years as the use of IaaS, PaaS, and SaaS solutions continues to increase.

Resource Policies

Key Topic

Resource policies involve creating, assigning, and managing rules over the cloud resources that systems (virtual machines, containers, and so on) or applications use. Resource policies are also often created to make sure that resources remain compliant with corporate standards and service-level agreements (SLAs). Cloud service providers offer different solutions (such as the Microsoft Azure Policy) that help you maintain good cloud usage governance by evaluating deployed resources and scanning for those not compliant with the policies that have been defined.

Let's look at a resource policy example. You could create a policy to allow only a certain size of virtual machines in the cloud environment. Once that policy is implemented, it would be enforced when you're creating and updating resources as well as over existing resources. Another example is a policy that can contain conditions and rules that determine if a storage account being deployed is within a certain set of criteria (size, how it is shared, and so on). Such a policy could deny all storage accounts that do not adhere to the defined set of rules.

Transit Gateway

Key Topic

A ***transit gateway*** is a system that can be used to interconnect a virtual private cloud (VPC) and on-premises networks. Transit gateways can also be used to connect cloud-hosted applications to software-defined wide-area network (SD-WAN) devices.

NOTE SD-WAN is a solution that simplifies the management of a wide-area network (WAN) implementation by decoupling network infrastructure devices from their control mechanism using software.

Figure 10-10 illustrates a transit gateway implementation. On-premises network infrastructure and systems are connected to the transit gateway using a VPN tunnel (typically using IPsec). The transit gateway then connects those resources to different VPCs.

Key Topic

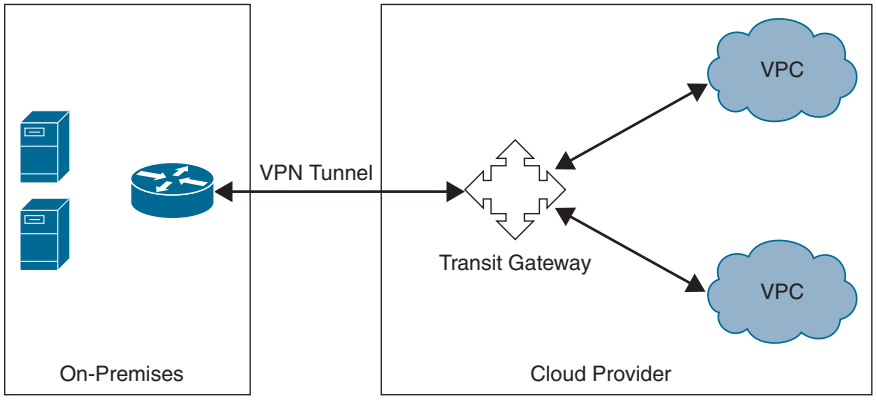


FIGURE 10-10 A Transit Gateway

Virtual Machine (VM) Sprawl Avoidance and VM Escape Protection

Virtualization is a solution that allows you to create IT services using hardware resources that can be shared among different installations of operating systems and applications. The individual virtual systems that run different operating systems and applications are called *virtual machines* (VMs). The hypervisor is the computer software that is used to create and run VMs. Organizations have used virtualization technologies for decades. There have been a few challenges in virtualized environments, however. One of those issues is VM sprawl. In the following sections, you learn the challenges of VM sprawl and how to potentially avoid it. You also learn about VM escape attacks and how to protect your organization against them.

Key Topic

Understanding and Avoiding VM Sprawl

Regardless of the size of your organization, you may find yourself in a situation in which too many virtual machines (or physical servers for that matter) may be deployed, which results in unnecessary complexity and inefficiency when those resources aren't used appropriately or are no longer needed. This unnecessary complexity also introduces security risks. ***VM sprawl*** (otherwise known as virtualization

sprawl) occurs when an organization can no longer effectively control and manage all the VMs on a network or in the cloud.

VM sprawl can happen with rapidly growing networks when numerous VMs are set up by different business units for dispersed applications—especially when you create temporary VMs that do not get purged. For instance, your developers or department might set up VMs as part of a testing environment; however, if the VMs aren't disposed of properly when they are no longer needed, your IT department can end up with too much to manage.

VM sprawl can impact your organization in a few different ways:

- Cost
- Management complexity and overhead
- Introduction of additional security risks
- Unnecessary disk space, CPU, and memory consumption
- Potential data protection challenges

One of the best ways to avoid VM sprawl (VM sprawl avoidance) is to have a robust, automated solution to perform audits of the VMs. This automated audit will allow you to identify which VMs are no longer needed and delete them. Additionally, you should have a good process to clean up orphaned VM snapshots and end unnecessary backups or replication of VMs no longer needed. It is also very important to use a good naming convention so that you can keep track of these VMs in your infrastructure or in the cloud.

Protecting Against VM Escape Attacks



In a virtualized environment, the hypervisor controls all aspects of VM operation. Securing a hypervisor is crucial but also fairly complex. There is an attack known as *hyper-jacking*. In this type of assault, an attacker or malware compromises one VM and then attacks the hypervisor. When a guest VM performs this type of attack, it is often called a ***VM escape attack***. In a VM escape attack, the guest VM breaks out of its isolated environment and attacks the hypervisor or could compromise other VMs hosted and controlled by the hypervisor. Figure 10-11 illustrates a VM escape attack.

Key Topic

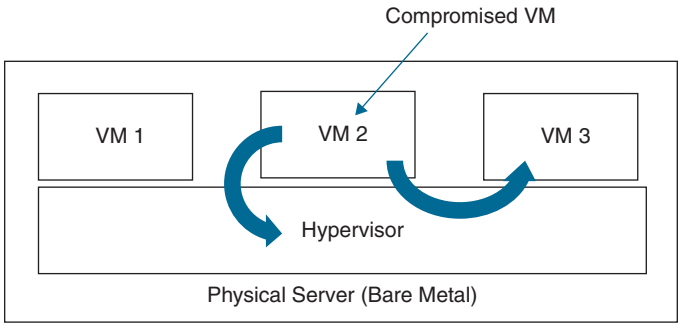


FIGURE 10-11 VM Escape Attack

Consider the following to mitigate the risk of VM escape attacks:

- Use a Type 1 hypervisor versus a Type 2 hypervisor to reduce the attack surface. A Type 1 hypervisor runs on physical servers or "bare metal," and Type 2 runs on top of an operating system. Type 1 hypervisors have a smaller footprint and reduced complexity.
- Harden the hypervisor's configuration per the recommendation of the vendor. For example, you can disable memory sharing between VMs running within the same hypervisor hosts.
- Disconnect unused external physical hardware devices (such as USB drives).
- Disable clipboard or file-sharing services.
- Perform self-integrity checks at boot-up using technologies like the Intel Trusted Platform Module/Trusted Execution Technology to verify whether the hypervisor has been compromised.
- Monitor and analyze hypervisor logs on an ongoing basis and look for indicators of compromise.
- Keep up with your hypervisor vendor's security advisories and apply security updates immediately after they are disclosed.
- Implement identity and access control across all tools, automated scripts, and applications using the hypervisor APIs.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 10-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 10-2 Key Topics for Chapter 10

Key Topic Element	Description	Page Number
List	Listing the cloud computing service categories	231
Tip	Surveying NIST's special publication (SP) 800-145	232
List	Listing the different types of clouds used by organizations	232
Paragraph	Defining edge computing	234
Figure 10-1	Edge Computing	235
Paragraph	Defining fog computing	235
Paragraph	Understanding what thin clients are	235
Figure 10-2	Thin Client Accessing a Virtualized Desktop	236
Paragraph	Understanding what containers are	236
Figure 10-3	How Container Images Work	237
Paragraph	Defining microservices	240
Figure 10-7	A High-Level SDN Implementation	242
Paragraph	Understanding cloud services integration	246
Paragraph	Defining resource policies	246
Paragraph	Understanding what a transit gateway is	246
Figure 10-10	A Transit Gateway	247
Section	Understanding and Avoiding VM Sprawl	247
Paragraph	Understanding what a VM escape attack is	248
Figure 10-11	VM Escape Attack	249

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

cloud computing, software as a service (SaaS), infrastructure as a service (IaaS), platform as a service (PaaS), anything as a service (XaaS), public cloud, private cloud, hybrid cloud, community cloud, cloud service provider (CSP), managed service provider (MSP), managed security service providers (MSSPs), off-premises, on-premises, edge computing, fog computing, thin clients, containers, microservices, application programming interfaces (APIs), software-defined networking (SDN), infrastructure as code, software-defined visibility (SDV), serverless architecture, services integration, resource policies, transit gateway, virtualization, VM sprawl, VM escape attack

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What system can be used to interconnect a virtual private cloud (VPC) and on-premises networks?
2. _____ are used in the process of creating, assigning, and managing rules over the cloud resources that systems (virtual machines, containers, and so on) or applications use.
3. What is a series of tools and technologies used to connect different systems, applications, code repositories, and physical or virtual network infrastructure to allow the real-time exchange of data and processes?
4. AWS Lambda is an example of which type of cloud service architecture?
5. OpenDaylight (ODL) is an example of a(n) _____ controller.



This chapter covers the following topics related to Objective 2.3 (Summarize secure application development, deployment, and automation concepts) of the CompTIA Security+ SY0-601 certification exam:

- Environment
  - Development
  - Test
  - Staging
  - Production
  - Quality assurance (QA)
- Provisioning and deprovisioning
- Integrity measurement
- Secure coding techniques
  - Normalization
  - Stored procedures
  - Obfuscation/camouflage
  - Code reuse/dead code
  - Server-side vs. client-side execution and validation
  - Memory management
  - Use of third-party libraries and software development kits (SDKs)
  - Data exposure
- Open Web Application Security Project (OWASP)
- Software diversity
  - Compiler
  - Binary
- Automation/scripting
  - Automated courses of action
  - Continuous monitoring
  - Continuous validation
  - Continuous integration
  - Continuous delivery
  - Continuous deployment
- Elasticity
- Scalability
- Version control

CHAPTER 11

Summarizing Secure Application Development, Deployment, and Automation Concepts

This chapter provides an overview of different software development environments and methodologies. You then learn how organizations provision and deprovision applications in modern environments. You also learn about software integrity measurement, different secure coding techniques, software diversity, and how automation and orchestration have a direct effect on elasticity and scalability.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table11-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 11-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Software Development Environments and Methodologies	1
Application Provisioning and Deprovisioning	2
Software Integrity Measurement	3
Secure Coding Techniques	4–6
Open Web Application Security Project (OWASP)	7
Software Diversity	8
Automation/Scripting	9
Elasticity and Scalability	10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is a software and hardware development and project management methodology that has at least five to seven phases that follow in strict linear order?
   a. Waterfall
   b. Agile
   c. DevOps
   d. SDLC

2. Which of the following is a benefit when you use log aggregation tools to maintain and analyze logs of every element that goes into the provisioning of applications?
   a. The ability to scale horizontally
   b. The ability to respond quickly and deprovision the application in the event that something wrong
   c. The ability to design an elastic infrastructure
   d. None of these answers are correct.

3. Which of the following elements can help software (code) integrity?
   a. Unit testing
   b. Integration testing
   c. Identifying a code integrity manager
   d. All of these answers are correct.

4. Which process includes identifying assets to the system or application, uncovering vulnerabilities, identifying threats, documenting threats, and rating those threats according to their potential impact?
   a. SECOPS
   b. Principle of least privilege
   c. Threat modeling
   d. None of these answers are correct.

5. Which of the following are important security principles that should be incorporated into the SDLC?
   a. Input validation
   b. Principle of least privilege
   c. Failing securely
   d. All of these answers are correct.

6. Which of the following might include syntax errors in the code and type-checking errors?
   a. Misconfigured VMs
   b. Unpatched applications and operating systems
   c. Misconfigured storage buckets
   d. Compile-time errors

7. Which of the following are top web application security risks?
   a. Broken Access Control
   b. XML External Entities (XXE)
   c. Cross-Site Scripting (XSS)
   d. All of these answers are correct.

8. Which of the following terms is often used when a compiler is modified to generate variants of a binary (target application) that operates in the same way when processing benign input but may operate in a different manner when given malicious input?
   a. ASLR
   b. Random forest
   c. Software diversity
   d. None of these answers are correct.

9. Which of the following is a software development practice where programmers merge code changes in a central repository multiple times a day?
   a. Continuous integration
   b. DevSecOps
   c. Waterfall
   d. All of these answers are correct.

10. Which of the following is the ability of an underlying infrastructure to react to a sudden increase in demand by provisioning more resources in an automated way?
    a. Load balancing
    b. Using Kubernetes
    c. Elasticity
    d. All of these answers are correct.

Foundation Topics

Software Development Environments and Methodologies

Before we start to define what DevOps is, let's look at the history of development methodologies and the underlying ***software development environments***. Decades of lessons learned from software development, high-reliability organizations, manufacturing, high-trust management models, and others have evolved into the DevOps practices used today.

The traditional software development methodology is the waterfall model, which is a software and hardware development and project management methodology that has at least five to seven phases that follow in strict linear order. Each phase cannot start until the previous phase has been completed. Figure 11-1 illustrates the typical phases of the waterfall development methodology.

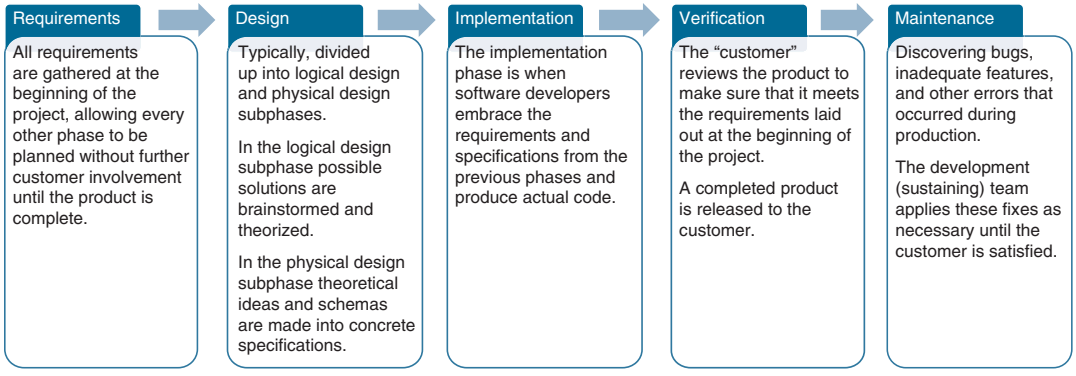


FIGURE 11-1 The Waterfall Development Methodology Phases

Organizations use the waterfall methodology for a few reasons. One of the main reasons is that project requirements are agreed upon from the beginning; consequently, planning and scheduling are simple and clear. With a fully laid-out project schedule, an accurate estimate can be given, including development project cost, resources, and deadlines. Another reason is that measuring progress is easy as the organization moves through the phases and hits the different milestones. End customers are not perpetually adding new requirements to the project, thus delaying production.

There also are several disadvantages to using the waterfall methodology. One of the disadvantages is that it can be difficult for customers to enumerate and communicate all of their needs at the beginning of the project. If end customers are dissatisfied with the product in the verification phase, going back and designing the code again

can be very costly. In the waterfall methodology, a linear project plan is rigid and lacks flexibility for adapting to unexpected events.

The next software development methodology is Agile. Agile is a software development and project management process in which a project is managed by breaking it up into several stages and involving constant collaboration with stakeholders and continuous improvement and iteration at every stage. The Agile methodology begins with end customers describing how the final product will be used and clearly articulating what problem it will solve. Once the coding begins, the respective team members cycle through a process of planning, executing, and evaluating. This process may allow the final deliverable to change in order to better fit the customers' needs. In an Agile environment, continuous collaboration is key. Clear and ongoing communication among team members and project stakeholders enables them to make fully informed decisions. The Agile methodology was originally developed in written form by 17 people in 2001, and it is documented at "The Manifesto for Agile Software Development" (https://agilemanifesto.org).

In Agile, the input to the development process is the creation of a business objective, concept, idea, or hypothesis. Then the work is added to a committed "backlog." From there, software development teams that follow the standard Agile or iterative process will transform that idea into "user stories" and some sort of feature specification. This specification is then implemented in code. The code is then checked in to a ***version control*** repository (for example, GitLab or GitHub), where each change is integrated and tested with the rest of the software system. In Agile, value is created only when services are running in production; consequently, the organization must ensure that it is not only delivering fast flow but also that deployments can be performed without causing chaos and disruptions, such as service outages, service impairments, or security or compliance failures. Figure 11-2 illustrates the general steps of the Agile methodology.

Many organizations adopt a concept related to Agile called Scrum. Scrum is a framework that helps organizations work together because it encourages teams to learn through experiences, self-organize while working on a solution, and reflect on their wins and losses to continuously improve. Scrum is used by software development teams; however, its principles and lessons can be applied to all kinds of teamwork. Scrum describes a set of meetings, tools, and roles that work in concert to help teams structure and manage their work.

TIP Scrum.org provides a set of resources, certification, and training materials related to Scrum.

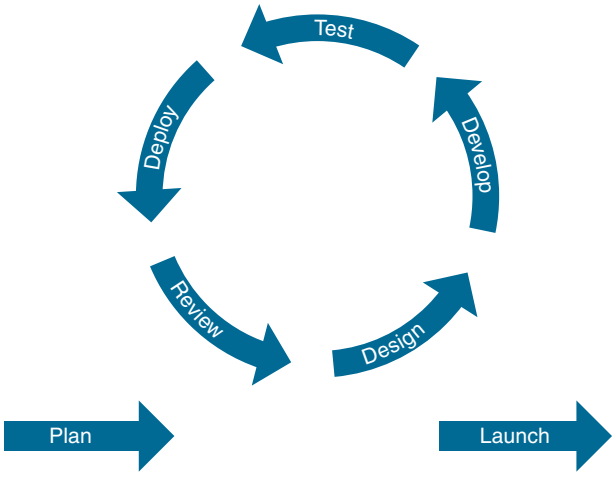


FIGURE 11-2 The Agile Methodology Phases

The increasing popularity of the Agile model led to DevOps and, subsequently, the practice of secure DevOps (or DevSecOps). DevOps is the outcome of many trusted principles—from software development, manufacturing, and leadership to the information technology value stream.

Regardless of the development methodology, most developers deal with different environments at some point. Each development environment has its own properties. The main environments are as follows:

Key Topic

- ***Development*:** In this environment you create your code on your computer or in the cloud. All of your commits and branches live here, along with those of your software development collaborators. Preliminary testing could happen in this environment. However, there is also a testing environment.
- ***Testing*:** This dedicated environment helps you make sure that your code does its job.
- ***Staging*:** This environment also allows you to test your code or application; however, it is as similar to the production environment as it can be. The staging environment enables you to ensure that each component of your application still does its job with everything else going on around it. Systems and your application in the staging environment connect to as many services as they can without affecting the production environment. A lot of the hardcore quality assurance testing takes place in the staging environment. Remember, staging environments reduce the risk of introducing issues before solutions are deployed in production.

- ***Production*:** This is the environment where your code fulfills its destiny (where users access the final code after all of the updates and testing). When you hear people talk about making "code go live," this is the environment they're talking about.
- ***Quality assurance (QA)*:** This is the process and methods of making sure that your organization and software developers create good quality software. A good QA process should include standards and procedures that managers, developers, and administrators could use to review and audit your software development.

Application Provisioning and Deprovisioning

Traditionally, configuring IT infrastructure, opening firewall ports, provisioning certificates, and performing other similar tasks fell in the hands of the network engineer. However, in modern environments, a lot of this provisioning happens automatically. In Chapter 10, "Summarizing Virtualization and Cloud Computing Concepts," you learned about the concept of provisioning and configuring the network and IT infrastructure as code (IaC). In the past, the network engineer's job was to log in to each device and do the necessary configurations manually, and the entire process was laborious and painstaking. Furthermore, ***application provisioning*** is typically not a one-person job. Application provisioning requires the coordination and cooperation of several stakeholders from multiple teams (that is, developers, application owners, IT network administrators, and security teams).

With the use of modern applications and tools like Ansible and Terraform, software developers and network engineers are creating and provisioning applications faster than ever. Using these solutions, you can provision an application in minutes rather than days or weeks.

Key Topic

Moreover, you can integrate with log aggregation tools to maintain and analyze logs of every element that goes into the provisioning. Therefore, you can respond quickly and ***deprovision*** the application in the event that something went wrong. Once you go back, you can check the logs and accurately find and fix the root cause of the error.

Traditionally, when an application needed to be fully retired or decommissioned, application deprovisioning was always a painful task. When you start the deprovisioning planning, you may not know how taking down one single application may have a direct effect on other applications that might have been integrated to it or that are using its data. Having a good automated process of logging, integration monitoring, and configuration will allow you to reduce the time needed to decommission the application and reduce the overall risk to the organization.

Software Integrity Measurement

Key Topic

Software integrity measurement (code integrity measurement) is the process of how to measure the source code's quality when it is passed on to the quality assessment (QA). This measurement is affected by how extensively the code was unit- and integration-tested. The purpose of code integrity measurement is for the developer to be sure that the code is written correctly. In short, code integrity checking helps organizations release better software with fewer bugs and security vulnerabilities.

Figure 11-3 includes a few elements that can help improve code integrity.

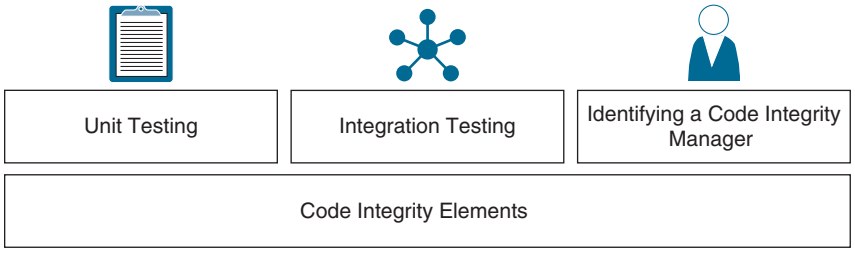


FIGURE 11-3 Elements That Help with Code Integrity

Figure 11-4 shows some of the advantages of code integrity checks and software integrity measurements.

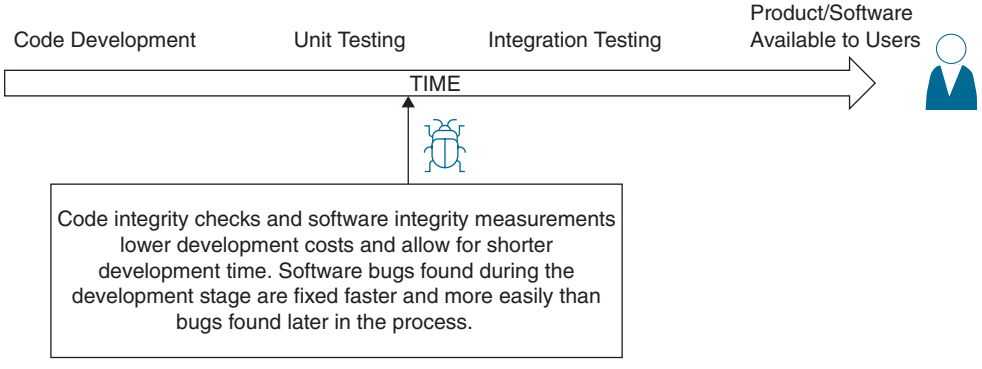


FIGURE 11-4 Advantages of Code Integrity Checks

Key Topic

Secure Coding Techniques

To properly develop a secure application, the developer has to scrutinize it at every turn, and from every angle, throughout the life of the project. Over time, this idea manifested itself into the concept known as the software development lifecycle

(SDLC)—an organized process of planning, developing, testing, deploying, and maintaining systems and applications, and the various methodologies used to do so. A common SDLC model used by companies is the waterfall model, mentioned previously. In this model, the SDLC is broken down into several sequential phases. Here's an example of an SDLC's phases based on the waterfall model:

1. **Planning and analysis.** Goals are determined, needs are assessed, and high-level planning is accomplished.
2. **Software/systems design.** The design of the system or application is defined and diagrammed in detail.
3. **Implementation.** The code for the project is written.
4. **Testing.** The system or application is checked thoroughly in a testing environment.
5. **Integration.** If multiple systems are involved, the application should be tested in conjunction with those systems.
6. **Deployment.** The system or application is put into production and is now available to end users.
7. **Maintenance.** Software is monitored and updated throughout the rest of its lifecycle. If there are many versions and configurations, version control is implemented to keep everything organized.

This is a basic example of the phases in an SDLC methodology. This one actually builds on the original waterfall model but is quite similar. However, it could be more or less complicated depending on the number of systems, the project goals, and the organization involved.

NOTE The SDLC is also referred to as the *software development process*, *systems development lifecycle*, and *application development lifecycle*. CompTIA uses the term *Software Development Life Cycle* in its Security+ objectives acronym list, but be ready for slightly different terms based on the same concept.

While variations of the waterfall model are commonplace, an organization might opt to use a different model, such as the V-shaped model, which stresses more testing, or rapid application development (RAD), which puts more emphasis on process and less emphasis on planning. As previously mentioned, the Agile model allows you to break the software development process into small increments and is designed to be more adaptive to change. It focuses on customer satisfaction, cooperation between developers and businesspeople, face-to-face conversation, simplicity, and quick adjustments to change.

The increasing popularity of the Agile model led to DevOps and, subsequently, the practice of secure DevOps. DevOps is a portmanteau of the terms *software development* and *information technology operations*. It emphasizes the collaboration of those two departments so that the coding, testing, and releasing of software can happen more efficiently—and hopefully, more securely. DevOps is similar to continuous delivery—which focuses on automation and quick execution—but from an organizational standpoint is broader and supports greater collaboration. A secure DevOps environment should include secure provisioning and deprovisioning of software, services, and infrastructure; security automation; continuous integration; baselining; infrastructure as code; and immutable systems. *Immutable* means unchanging over time. From a systems and infrastructure viewpoint, it means that software and services are replaced instead of being changed. Rapid, efficient deployment of new applications is at the core of DevOps, and it is one of the main tasks of all groups involved.

Core SDLC and DevOps Principles

The terms and models discussed in the previous section can be confusing at times because they are similar and because there are overlapping gray areas between them. So, if you are ever confused while working in the field, it's good to think back to the core fundamentals of security. From a larger perspective, you, as security administrator, and the programmer/systems developer should always keep the CIA concept in mind:

- **Maintaining confidentiality:** Allowing users access only to data to which they have permission
- **Preserving integrity:** Ensuring data is not tampered with or altered
- **Protecting availability:** Ensuring systems and data are accessible to authorized users when necessary

The CIA concepts are important when doing a secure code review, which can be defined as an in-depth code inspection procedure. Organizations often include it as part of the testing phase of the SDLC, but it is usually conducted before other tests such as fuzzing or penetration tests, which are discussed later in this chapter.

In general, quality assurance policies and procedures should be implemented while developing and testing code. This implementation will vary from one organization to the next but generally includes procedures that have been developed by a team, a set of checks and balances, and a large amount of documentation. By checking the documentation within a project, a developer can save a lot of time, while keeping the project more secure.

From a larger perspective, an organization might implement modeling as part of its software quality assurance program. By modeling, or simulating, a system or application, the organization can test, verify, validate, and finally accredit it as acceptable for a specific purpose.

One secure and structured approach that organizations take is called threat modeling. Threat modeling enables you to prioritize threats to an application based on their potential impact. This modeling process includes identifying assets to the system or application, uncovering vulnerabilities, identifying threats, documenting threats, and rating those threats according to their potential impact. The more risk, the higher the rating. Threat modeling is often incorporated into the SDLC during the design, testing, and deployment phases.

Some other important security principles that should be incorporated into the SDLC include

- **Apply the principle of least privilege:** Applications should be coded and run in such a way as to maintain the principle of least privilege. Users should have access only to what they need. Processes should run with only the bare minimum access needed to complete their functions. However, this can be coupled with separation of privilege, where access to objects depends on more than one condition (for example, authentication plus an encrypted key).
- **Apply the principle of defense in depth:** The more security controls, the better. The layering of defense in secure coding may take the form of validation, encryption, auditing, special authentication techniques, and so on.
- **Ensure applications never trust user input:** Input should be validated carefully using input validation techniques.
- **Minimize the attack surface area:** Every additional feature that a programmer adds to an application increases the size of the attack surface and increases risk. Unnecessary functions should be removed, and necessary functions should require authorization.
- **Establish secure defaults:** Out-of-the-box offerings should be as secure as possible. If possible, user password complexity and password aging default policies should be configured by the programmer, not the user. Permissions should default to no access and should be granted only as they are needed.
- **Provide for authenticity and integrity:** For example, when deploying applications and scripts, use *code signing* in the form of a cryptographic hash with verifiable checksum for validation. A digital signature will verify the author and/or the version of the code—that is, if the corresponding private key is secured properly.

- **Fail securely:** At times, applications will fail. How they fail determines their security. Failure exceptions might show the programming language that was used to build the application, or worse, lead to access holes. Error handling/exception handling code should be checked thoroughly so that a malicious user can't find out any additional information about the system. These error-handling methods are sometimes referred to technically as pseudocodes. For example, to handle a program exception, a properly written pseudocode will basically state (in spoken English): "If a program module crashes, then restart the program module."
- **Fix security issues correctly:** Once found, security vulnerabilities should be thoroughly tested, documented, and understood. Patches should be developed to fix the problem, but not cause other issues or application regression.

Key Topic

There are some other concepts to consider. First is ***obfuscation*** (or ***camouflage***), which is the complicating of source code to make it more difficult for people to understand. Code is obfuscated to conceal its purpose in order to prevent tampering and/or reverse engineering. It is an example of security through obscurity. Other examples of security through obscurity include *code camouflaging* and *steganography*, both of which might be used to hide the true code being used. Another important concept is *code checking*, which involves limiting the reuse of code to that which has been approved for use, and removing dead code. It's also vital to incorporate good ***memory management*** techniques. Secure memory management and memory protection techniques include different ways to control memory access rights on a system to prevent a process from accessing memory that has not been allocated to it. Memory management best practices prevent memory-related vulnerabilities and potential malware within a process from affecting other processes or the underlying operating system. Secure memory protection should include all accesses to a specified area of memory, write accesses, or attempts to execute the contents of the memory area. Examples of memory protection techniques include address space layout randomization (ASLR) and executable space protection.

Finally, be very careful when using ***third-party libraries and software development kits (SDKs)***, and test them thoroughly before using them within a live application.

For the Security+ exam, the most important of the SDLC phases are maintenance and testing. In the maintenance phase, which doesn't end until the software is removed from all computers, an application needs to be updated accordingly, corrected when it fails, and constantly monitored. In the testing phase, a programmer (or team of programmers and other employees) checks for bugs and errors in a variety of ways. It's imperative that you know some of the vulnerabilities and attacks to a system or application and also how to fix them and protect against them. The best way to prevent these attacks is to test and review code.

Programming Testing Methods

Several testing methods and techniques can be implemented to seek out programming vulnerabilities and help prevent attacks from happening. Programmers have various ways to test their code, including system testing, input validation, and fuzzing. By using a combination of these testing techniques during the testing phase of the SDLC, there is a much higher probability of a secure application as the end result.

In some security assessments and testing, knowledge of the system code—and programming knowledge in general—is not required. The tester does not know about the system's internal structure and is often given limited information about what the application or system is supposed to do. In this type of testing, one of the most common goals is to crash the program. If a user is able to crash the program by entering improper input, the programmer has probably neglected to thoroughly check the error-handling code and/or input validation.

On the other side of the spectrum, testers have in-depth architectural knowledge and are given detailed information about the design of the system. They are given login details, production documentation, and source code. System testers might use a combination of fuzzing (covered shortly), data flow testing, and other techniques such as stress testing, penetration testing, and sandboxes. *Stress testing* is usually done on real-time operating systems, mission-critical systems, and software, and checks whether they have the robustness and availability required by the organization. As you learned in Chapter 8, "Understanding the Techniques Used in Penetration Testing," a *penetration test* is a method of evaluating a system's security by simulating one or more attacks on that system. *Sandbox* is the term applied when a web script (or other code) runs in its own environment (often a virtual environment) for the express purpose of not interfering with other processes, often for testing. Sandboxing technology is frequently used to test unverified applications for malware, malignant code, and possible errors such as buffer overflows.

Compile-Time Errors vs. Runtime Errors

Programmers and developers need to test for potential compile-time errors and runtime errors. Compile time refers to the duration of time during which the statements written in any programming language are checked for errors. *Compile-time errors* might include syntax errors in the code and type-checking errors. A programmer can check these without actually running the program and instead check it in the compile stage when it is converted into machine code.

A *runtime error* is a program error that occurs while the program is running. The term is often used in contrast to other types of program errors, such as syntax errors and compile-time errors. Runtime errors might include running out of memory, invalid memory address access, invalid parameter value, or buffer overflows/dereferencing a null pointer (to name a few), all of which can be discovered only by running the program as a user. Another potential runtime error can occur if there is an attempt to divide by zero. These types of errors result in a software exception. Software and hardware exceptions need to be handled properly. Consequently, structured exception handling (SEH) is a mechanism used to handle both types of exceptions. It enables the programmer to have complete control over how exceptions are handled and provides support for debugging.

Code issues and errors that occur at either compile time or run time could lead to vulnerabilities in the software. However, it's the runtime environment that we are more interested in from a security perspective because that more often is where attackers will attempt to exploit software and websites.

Input Validation

Input validation is very important for website design and application development. *Input validation*, or data validation, is a process that ensures the correct usage of data: it checks the data that is input by users into web forms and other similar web elements. If data is not validated correctly, it can lead to various security vulnerabilities including sensitive ***data exposure*** and the possibility of data corruption. You can validate data in many ways, from coded data checks and consistency checks to spelling and grammar checks, and so on.

Key Topic

Attackers can perform ***server-side execution*** by leveraging vulnerabilities such as command injection, cross-site request forgery, and SQL injection. In other words, attackers can manipulate an application by providing crafted input to an application to perform a malicious transaction, steal sensitive data, or cause a denial-of-service (DoS) condition. Similarly, attackers can perform ***client-side execution*** by exploiting vulnerabilities that can perform malicious transactions in a user's endpoint, web browser, or mobile device. All application programming interfaces (APIs) and language functions that accept data can be potentially sent malicious input to cause the device or client application to crash or—what is worse—perform code execution.

Whatever data is being dealt with, it should be checked to make sure it is being entered correctly and won't create or take advantage of a security flaw. If validated properly, bad data and malformed data will be rejected. Input validation should be done both on the client side and, more importantly, on the server side.

If an organization has a web page with a PHP-based contact form, for example, the data entered by the visitor should be checked for errors or maliciously typed input. The following piece of PHP code is contained within a common contact form:

```
else if (!preg_match('/^[A-Za-z0-9.-]+$/', $domain))
 {
 // character not valid in domain part
 $isValid = false;
 }
```

This snippet is part of a larger piece of code that checks the entire email address a user enters into a form field. This particular code snippet checks to make sure the user is not trying to enter a backslash in the domain name portion of the email address. A backslash is not allowed in email addresses and could be detrimental if used maliciously. In the first line within the brackets, the code says **A-Za-z0-9.-**, which tells the system what characters are allowed. Uppercase and lowercase letters, numbers, periods, and dashes are allowed, but other characters such as backslashes, dollar signs, and so on are not allowed. Those other characters would be interpreted by the form's supporting PHP files as illegitimate data and would not be passed on through the system. The user would receive an error, which is a part of ***client-side validation***. But the more a PHP form is programmed to check for errors, the more it is possible to have additional security holes. Therefore, ***server-side validation*** is even more important. Any data that *is* passed on by the PHP form should be checked at the server as well. In fact, an attacker might not even be using the form in question but might be attacking the URL of the web page in some other manner. This type of attack can be checked at the server within the database software or through other means. By the way, the concept of combining client-side and server-side validation also goes for pages that utilize JavaScript.

This is just one basic example, but as mentioned previously, input validation is the key to preventing attacks such as SQL injection and XSS. All form fields should be tested for good input validation code, both on the client side and the server side. By combining the two and checking every access attempt, you develop complete mediation of requests.

TIP Using input validation is one way to prevent sensitive data exposure, which occurs when an application does not adequately protect personally identifiable information (PII). Such exposure can also be prevented by making sure that input data is never stored or transmitted in clear text; using strong encryption and securing key generation and storage; using HTTPS for authentication; and using a salt, which is random data used to strengthen hashed passwords.

Static and Dynamic Code Analysis

Static code analysis is a type of debugging that is carried out by examining the code *without* executing the program. This analysis can be done by scrutinizing code visually or with the aid of specific automated tools—static code analyzers—based on the language being used. Static code analysis can help reveal major issues with code that could even lead to disasters. Although this phase of testing is important, it should always be followed by some type of dynamic analysis—for example, fuzz testing (covered next). At this point, the program is actually executed while it is being tested. It is meant to locate minor defects in code and vulnerabilities.

Fuzz Testing

Fuzz testing (also known as fuzzing or dynamic analysis) is another smart concept. In this type of analysis, random data is input into a computer program in an attempt to find vulnerabilities. This is often done without knowledge of the source code of the program. The program to be tested is run, has data input to it, and is monitored for exceptions such as crashes. This analysis can be done with applications and operating systems. It is commonly used to check file parsers within applications such as Microsoft Word and network parsers that are used by protocols such as DHCP.

Fuzz testing, or fuzzing, can be used to find software errors (or bugs) and security vulnerabilities in applications, operating systems, infrastructure devices, IoT devices, and other computing devices. Fuzzing involves sending random data to the unit being tested in order to find input validation issues, program failures, buffer overflows, and other flaws. Tools that are used to perform fuzzing are referred to as "fuzzers." Examples of popular fuzzers are Peach, Munity, American Fuzzy Lop, and Synopsys Defensics.

The Mutiny Fuzzing Framework is an open-source fuzzer created by Cisco. It works by replaying packet capture files (PCAPs) through a mutational fuzzer. You can obtain more information about Mutiny Fuzzing Framework at https://github.com/Cisco-Talos/mutiny-fuzzer. The Mutiny Fuzzing Framework uses Radamsa to perform mutations. Radamsa is a tool that can be used to generate test cases for fuzzers. You can download additional information about Radamsa at https://gitlab.com/akihe/radamsa.

American Fuzzy Lop (AFL) is a tool that provides features of compile-time instrumentation and genetic algorithms to automatically improve the functional coverage of fuzzing test cases. You can obtain additional information about AFL at http://lcamtuf.coredump.cx/afl/.

Peach is one of the most popular fuzzers in the industry. There is a free (open-source) version, the Peach Fuzzer Community Edition, and a commercial version. You can download the Peach Fuzzer Community Edition and obtain additional information about the commercial version at https://gitlab.com/gitlab-org/security-products/protocol-fuzzer-ce.

TIP For additional examples of fuzzers and fuzzing, see the GitHub repository at https://h4cker.org/github.

Fuzz testing can uncover full system failures, memory leaks, and error-handling issues. Fuzzing is usually automated (a program that checks a program) and can be as simple as sending a random set of bits to be input to the software. However, designing the inputs that cause the software to fail can be a tricky business, and often a myriad of variations of code must be tried to find vulnerabilities. Once the fuzz test is complete, the results are analyzed, the code is reviewed and made stronger, and vulnerabilities that were found are removed. The stronger the fuzz test, the better the chances that the program will not be susceptible to exploits.

Key Topic

As a final word on this type of analysis, after code is properly tested and approved, it should be reused whenever possible. This way, you can avoid reinventing the wheel and prevent common mistakes that a programmer might make that others may have already fixed. Just remember, ***code reuse*** is applicable only if the code is up to date and approved for use. Code reuse and ***dead code*** can introduce security risks if the code that is being shared is not properly tracked and patched. For instance, you may be patching vulnerabilities in one product or application that uses the original code, but not in others that reuse that code.

Programming Vulnerabilities and Attacks

The following sections address program code vulnerabilities and the attacks that exploit them. They also give specific ways to mitigate these risks during application development, hopefully preventing these threats and attacks from becoming realities.

NOTE This relatively short discussion covers a topic that could fill several books. It is not a seminar on how to program, but rather a concise description of programmed attack methods and some defenses against them. The Security+ objectives do not go into great detail concerning these methods, but CompTIA does expect you to have a broad and basic understanding.

Testing for Backdoors

Applications should be analyzed for backdoors. Backdoors are used in computer programs to bypass normal authentication and other security mechanisms in place. They can be avoided by updating the operating system and applications and firmware on devices, and especially by carefully checking the code of the program. If the system is not updated, a malicious person could take all kinds of actions via the backdoor. For example, a company's software developer could install code through a backdoor that reactivates the user account after it is disabled (for whatever reason—termination, end of consulting period, and so on). This is done through the use of a logic bomb (in addition to the backdoor) and could be deadly when the software developer has access again. To reiterate, you should make sure the operating system is updated and also consider job rotation, where programmers check each other's work.

Memory/Buffer Vulnerabilities

As you learned in Chapter 3, "Analyzing Potential Indicators Associated with Application Attacks," *memory management* and buffer vulnerabilities have been used by attackers to compromise systems in different ways. Of the several types, perhaps most important is the buffer overflow. A buffer overflow occurs when a process stores data outside the memory that the developer intended. This overflow could cause erratic behavior in the application, especially if the memory already had other data in it.

TIP Stacks and heaps are data structures that can be affected by buffer overflows. The stack is a key data structure necessary for the exchange of data between procedures. The heap contains data items whose size can be altered during execution.

In Chapter 3 you also learned what *integer overflows* and *memory leaks* are.

Another potential memory-related issue deals with pointer dereferencing—for example, the null pointer dereference. Pointer dereferencing is common in programming; when you want to access data (say, an integer) in memory, dereferencing the pointer would retrieve different data from a different section of memory (perhaps a different integer). Programs that contain a *null pointer dereference* generate memory fault errors (memory leaks). A null pointer dereference occurs when the program dereferences a pointer that it expects to be valid but is null, which can cause the application to exit or the system to crash. From a programmatical standpoint, the main way to prevent this issue is meticulous coding. Programmers can use special memory error analysis tools to enable error detection for a null pointer dereference. Once the problem is identified, the programmer can correct the code that may be causing these errors. But this concept can be used to attack systems over

the network by initiating IP address to hostname resolutions—ones that the attacker hopes will fail—causing a return null. What this all means is that the network needs to be protected from attackers attempting this (and many other) programmatical and memory-based attacks via a network connection.

A programmer may make use of address space layout randomization (ASLR) to help prevent the exploitation of memory corruption vulnerabilities. It randomly arranges the different address spaces used by a program (or process). This can aid in protecting mobile devices (and other systems) from exploits caused by memory-management problems. While there are exploits to ASLR (such as side-channel attacks) that can bypass it and derandomize how the address space is arranged, many systems employ some version of ASLR.

XSS and XSRF

Two web application vulnerabilities to watch out for include *cross-site scripting (XSS)* and *cross-site request forgery (XSRF)*.

XSS holes are vulnerabilities that can be exploited with a type of code injection. Code injection is the exploitation of a computer programming bug or flaw by inserting and processing invalid information; it is used to change how the program executes data. In the case of an XSS attack, an attacker inserts malicious scripts into a web page in the hope of gaining elevated privileges and access to session cookies and other information stored by a user's web browser. This code (often JavaScript) is usually injected from a separate "attack site." It can also manifest itself as an embedded JavaScript image tag, header manipulation (as in manipulated HTTP response headers), or other HTML-embedded image object within emails (that are web-based). Programmers can defeat the XSS attack through the use of output encoding (JavaScript escaping, CSS escaping, and URL encoding), by preventing the use of HTML tags, and by input validation: for example, checking forms and confirming that input from users does not contain hypertext. On the user side, the possibility of this attack's success can be reduced by increasing cookie security and by disabling scripts. If XSS attacks by email are a concern, the user could opt to set the email client to text only.

The XSS attack exploits the trust a user's browser has in a website. The converse of this, the XSRF attack, exploits the trust that a website has in a user's browser. In this attack (also known as a one-click attack), the user's browser is compromised and transmits unauthorized commands to the website. The chances of this attack can be reduced by requiring tokens on web pages that contain forms, using special authentication techniques (possibly encrypted), scanning .XML files (which could contain the code required for unauthorized access), and submitting cookies twice instead of once, while verifying that both cookie submissions match.

More Code Injection Examples

Other examples of code injection include SQL injection, XML injection, and Lightweight Directory Access Protocol (LDAP) injection.

Databases are just as vulnerable as web servers. The most common kind of database is the relational database, which is administered by a relational database management system (RDBMS). These systems are usually written in the Structured Query Language (SQL, pronounced "sequel"). An example of a SQL database is Microsoft's SQL Server; it can act as the back end for a program written in Visual Basic or Visual C++. Another example is MySQL, a free, open-source relational database often used in conjunction with websites that employ PHP pages.

One concern with SQL is the SQL injection attack, which occurs in databases, ASP.NET applications, and blogging software (such as WordPress) that use MySQL as a back end. In these attacks, user input in web forms is not filtered correctly and is executed improperly, with the end result of gaining access to resources or changing data. For example, the login form for a web page that uses a SQL back end (such as a WordPress login page) can be insecure, especially if the front-end application is not updated. An attacker will attempt to access the database (from a form or in a variety of other ways), query the database, find a user, and then inject code to the password portion of the SQL code—perhaps something as simple as **X = X**. This will allow any password for the user account to be used. If the login script was written properly (and validated properly), it should deflect this injected code. But if not, or if the application being used is not updated, it could be susceptible. An attack can be defended against by constraining user input, filtering user input, and using stored procedures such as input validating forms. Used to save memory, a ***stored procedure*** is a subroutine in an RDBMS that is typically implemented as a data-validation or access-control mechanism which includes several SQL statements in one procedure that can be accessed by multiple applications.

Key Topic

When using relational databases such as SQL and MySQL, a programmer works with the concept of ***normalization***, which is the ability to avoid or reduce data redundancies and anomalies—a core concept within relational DBs. There are, however, other databases that work within the principle of denormalization and don't use SQL (or use code in addition to SQL). Known as NoSQL databases, they offer a different mechanism for retrieving data than their relational database counterparts. These databases are commonly found in data warehouses and virtual systems provided by cloud-based services. While they are usually resistant to SQL injection, there are NoSQL injection attacks as well. Because of the type of programming used in NoSQL, the potential impact of a NoSQL injection attack can be greater than that of a SQL injection attack. An example of a NoSQL injection attack is the JavaScript Object Notation (JSON) injection attack. But NoSQL databases are also vulnerable to brute-force attacks (cracking of passwords) and connection

pollution (a combination of XSS and code injection techniques). Methods to protect against NoSQL injection are similar to the methods mentioned for SQL injection. However, because NoSQL databases are often used within cloud services, you, as security administrator, might not have much control over the level of security that is implemented. In these cases, careful scrutiny of the service-level agreement (SLA) between the company and the cloud provider is imperative.

LDAP injection is similar to SQL injection, again using a web form input box to gain access or by exploiting weak LDAP lookup configurations. The Lightweight Directory Access Protocol is used to maintain a directory of information such as user accounts or other types of objects. The best way to protect against this (and all code injection techniques for that matter) is to incorporate strong input validation.

XML injection attacks can compromise the logic of Extensible Markup Language applications—for example, XML structures that contain the code for users. They can be used to create new users and possibly obtain administrative access. You can test for these attacks by attempting to insert XML metacharacters such as single and double quotes. They can be prevented by filtering *in* allowed characters (for example, A–Z only). This is an example of "default deny," where only what you explicitly filter in is permitted; everything else is forbidden.

One point to remember is that when attackers utilize code injecting techniques, they are adding their own code to existing code or are inserting their own code into a form. A variant of this technique is command injection, which doesn't utilize new code; instead, an attacker executes system-level commands on a vulnerable application or OS. The attacker might enter the command (and other syntax) into an HTML form field or other web-based form to gain access to a web server's password files.

NOTE Though not completely related, another type of injection attack is DLL injection. In this attack, code is run within the address space of another process by forcing it to load a dynamic link library (DLL). Ultimately, this attack can influence the behavior of a program that was not originally intended.

Once again, the best way to defend against code injection/command injection techniques in general is to implement input validation during the development, testing, and maintenance phases of the SDLC.

Directory Traversal

Directory traversal, or the ../ (dot-dot-slash) attack, is a method of accessing unauthorized parent (or worse, root) directories. It is often used on web servers that have PHP files and are Linux or UNIX-based, but it can also be perpetrated on

Microsoft operating systems (in which case, it would be ..\ or the dot-dot-backslash attack). It is designed to get access to files such as ones that contain passwords. This access can be prevented by updating the operating system or by checking the code of files for vulnerabilities, otherwise known as fuzzing. For example, a PHP file on a Linux-based web server might have a vulnerable **if** or **include** statement, which when attacked properly could give the attacker access to higher directories and the passwd file.

Zero-Day Attack

A *zero-day attack* is executed on a vulnerability in software before that vulnerability is known to the creator of the software and before the developer can create a patch to fix the vulnerability. It's not a specific attack, but rather a group of attacks including viruses, Trojans, buffer overflow attacks, and so on. These attacks can cause damage even after the creator knows of the vulnerability because time may be needed to release a patch to prevent the attacks and fix damage caused by them. These attacks can be discovered by thorough analysis and fuzz testing.

Zero-day attacks can be prevented by using newer operating systems that have protection mechanisms and by updating those operating systems. They can also be prevented by using multiple layers of firewalls and by using whitelisting, which allows only known good applications to run. Collectively, these preventive methods are referred to as zero-day protection.

Table 11-2 summarizes the programming vulnerabilities/attacks covered in this section.

Table 11-2 Summary of Programming Vulnerabilities and Attacks

Vulnerability	Description
Backdoor	An attack placed by programmers, knowingly or inadvertently, to bypass normal authentication and other security mechanisms in place.
Buffer overflow	A vulnerability that occurs when a process stores data outside the memory that the developer intended.
Remote code execution (RCE)	An attack that occurs when an attacker obtains control of a target computer through some sort of vulnerability, gaining the power to execute commands on that remote computer.
Cross-site scripting (XSS)	An attack that exploits the trust a user's browser has in a website through code injection, often in web forms.
Cross-site request forgery (XSRF)	An attack that exploits the trust that a website has in a user's browser, which becomes compromised and transmits unauthorized commands to the website.

Vulnerability	Description
Code injection	An attack that occurs when user input in database web forms is not filtered correctly and is executed improperly. SQL injection is a very common example.
Directory traversal	A method of accessing unauthorized parent (or root) directories.
Zero day	A group of attacks executed on vulnerabilities in software before those vulnerabilities are known to the creator.

The CompTIA Security+ exam objectives don't expect you to be a programmer, but they do expect you to have a basic knowledge of programming languages and methodologies so that you can help secure applications effectively. I recommend a basic knowledge of programming languages used to build applications, such as Visual Basic, C++, C#, Java, and Python, as well as web-based programming languages such as HTML, ASP, and PHP, plus knowledge of database programming languages such as SQL. This foundational knowledge will help you not only on the exam but also when you, as security administrator, have to act as a liaison to the programming team or if you are actually involved in testing an application.



Open Web Application Security Project (OWASP)

A great resource mentioned numerous times in this book is the ***Open Web Application Security Project (OWASP)***, accessed at https://owasp.org. OWASP is a nonprofit organization that has chapters all over the world that focus on application and software security. It has numerous well-known and comprehensive projects designed to increase the awareness of secure coding and testing, and to create tools to help find and prevent security vulnerabilities.

The OWASP Testing Project is a comprehensive guide focused on web application testing. It is a compilation of many years' worth of work by OWASP members. It covers the high-level phases of web application security testing and also digs deeper into the actual testing methods used. For instance, it goes as far as providing strings for testing cross-site scripting (XSS) and SQL injection attacks. From a web application security testing perspective, it is the most detailed and comprehensive guide available. The OWASP Testing Project is available at www.owasp.org/index.php/OWASP_Testing_Project.

The OWASP Proactive Controls (www.owasp.org/index.php/OWASP_Proactive_Controls) is a collection of secure development practices and guidelines that any

software developer should follow to build secure applications. These practices will help you shift security earlier into design, coding, and testing. Here are the OWASP Top 10 Proactive Controls:

1. Define Security Requirements
2. Leverage Security Frameworks and Libraries
3. Secure Database Access
4. Encode and Escape Data
5. Validate All Inputs
6. Implement Digital Identity
7. Enforce Access Controls
8. Protect Data Everywhere
9. Implement Security Logging and Monitoring
10. Handle All Errors and Exceptions

One of the most popular OWASP projects is the Top 10 Web Application Security Risks. You can find the latest Top 10 Web Application Security Risks at https://owasp.org/www-project-top-ten. The following are the OWASP Top 10 Web Application Security Risks at the time of writing:

1. Injection
2. Broken Authentication
3. Sensitive Data Exposure
4. XML External Entities (XXE)
5. Broken Access Control
6. Security Misconfiguration
7. Cross-Site Scripting (XSS)
8. Insecure Deserialization
9. Using Components with Known Vulnerabilities
10. Insufficient Logging and Monitoring

Software Diversity

Typically, organizations have different system variants that are developed or acquired simultaneously to address a wide range of business or customer requirements. This variation is referred to as ***software diversity***. Software diversity has a direct impact on all phases of software development. In some cases, this diversity could lead to an increase of complexity because this variance and diversity has to be managed in requirements analysis, design, implementation, and validation. However, in some cases, there are different benefits of having software diversity within your environment.

For example, applications or systems providing the same service but that are from different vendors or developed by different groups inside the corporation may not rely on a single operating system or architecture. In some cases, they either do not have the same vulnerability or cannot be compromised with the same exploit. Some organizations claim that the use of software diversity increases attack tolerance. Others claim that it increases complexity in the patch management process.

There is another aspect of software diversity. This is when a ***compiler*** is modified to generate variants of a ***binary*** (target application) that operate in the same way when processing benign input; however, it might operate in a different manner when given malicious input. This situation is different from the previous examples of using a diverse set of applications and technologies. This new aspect of software diversity is handled by generating variants of a program by building a binary with a diversifying compiler that can randomize the code layout, stack variables, and random allocations of heap objects at different locations in each variant.

Automation/Scripting

Automation and ***scripting*** are crucial nowadays. Companies are finding new ways to automate everything (including software development, operations, and security). Earlier in this chapter you learned about DevOps, and subsequently, the practice of secure DevOps (or DevSecOps).

A secure DevOps environment should include the following concepts:

Key Topic

- ***Automated courses of action*:** These courses of action ensure secure provisioning and deprovisioning of software, services, and infrastructure.
- ***Continuous monitoring*:** Monitoring ensures that applications and systems are operating correctly and securely.
- ***Continuous validation*:** This process allows you to validate applications and code in an automated fashion.

- ***Continuous integration (CI)* and *continuous delivery (CD)*:** CI is a software development practice in which programmers merge code changes in a central repository multiple times a day. CD sits on top of CI and provides a way for automating the entire software release process. When you adopt CI/CD methodologies, each change in code should trigger an automated build-and-test sequence. This automation should also provide feedback to the programmers who made the change. CI/CD has been adopted by many organizations that provide cloud services (that is, SaaS, PaaS, and so on). For instance, CD can include cloud infrastructure provisioning and deployment, which traditionally have been done manually and consist of multiple stages. The main goal of the CI/CD processes is to be fully automated, with each run fully logged and visible to the entire team.
- ***Continuous deployment*:** This concept ensures the automation of the application deployment, provisioning, and underlying network components and infrastructure.
- ***Version control*:** In modern environments, code is checked in to a version control repository (for example, GitLab or GitHub), where each change is integrated and tested with the rest of the software system. Organizations that do not have proper version control will have a hard time keeping track of bug fixes and security patches. Similarly, vendors and software providers that do not employ appropriate version control and tracking make it harder for the consumers of their technology to be able to correlate, triage, and patch security vulnerability fixes.

Key Topic

Elasticity and Scalability

Many organizations move to the cloud because public cloud providers have the ability to scale their services indefinitely to provide unlimited computing, storing, and networking functionality. No environment is immune to failure or ***scalability*** problems; however, large cloud providers can scale due to the gigantic size of their shared resource pools. An *elastic* infrastructure can provide a scalable solution. ***Elasticity*** is the ability of an underlying infrastructure to react to a sudden increase in demand by provisioning more resources in an automated way.

Elasticity and scalability are often achieved by deploying technologies such as load balancers and by deploying applications and resources in multiple geographical locations (data centers around the world). Other technologies, such as enabling concurrent processing (parallel processing) and automated container deployments (such as using Kubernetes) allow organizations to auto-scale. To auto-scale, systems need to be adaptive enough to perform replication of state, fail-over, and upgrades without any manual instruction, intervention, and interpretation. Additionally, modern

implementations often use software infrastructure solutions including API gateways and service mesh solutions to ensure the overall system resiliency. In short, automation and orchestration are the key requirements for a reliable, scalable, and elastic IT infrastructure.

Earlier in this chapter, you learned that DevOps highlights the collaboration of software developers and IT infrastructure engineers so that the coding, testing, and releasing of software can happen more efficiently and more securely. DevOps concepts were created to shorten the software development lifecycle and provide continuous delivery through quality software. The term *elasticity* describes the capability to scale up and scale down applications. A good "elastic" environment is designed to make sure that resources are allocated appropriately and only when they are needed. With resource pooling, resources are brought together in a pooled model to provide services to many customers.

As you previously learned in this chapter, infrastructure as code allows for infrastructure configuration to be incorporated into application code, thus enabling DevOps teams to test applications in production-like environments from the beginning of the development cycle.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 11-3 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 11-3 Key Topics for Chapter 11

Key Topic Element	Description	Page Number
List	Identifying different development environments	259
Paragraph	Understanding software deprovisioning	260
Paragraph	Defining software integrity measurement	261
Section	Secure Coding Techniques	261

Key Topic Element	Description	Page Number
Paragraph	Understanding obfuscation/camouflage, memory management, and the risk associated with different third-party libraries and software development kits (SDKs)	265
Paragraph	Understanding server-side vs. client-side execution	267
Paragraph	Understanding the risk associated with code reuse and dead code	270
Paragraph	Defining the concept of normalization	273
Section	Open Web Application Security Project (OWASP)	276
List	Understanding different elements of the secure DevOps environment	278
Section	Elasticity and Scalability	279

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

software development environments, version control, development, testing, staging, production, quality assurance (QA), application provisioning, deprovision, software integrity measurement, obfuscation, camouflage, memory management, third-party libraries and software development kits (SDKs), data exposure, server-side execution, client-side execution, client-side validation, server-side validation, code reuse, dead code, stored procedure, normalization, Open Web Application Security Project (OWASP), software diversity, compiler, binary, automation, scripting, continuous delivery, continuous integration, automated courses of action, continuous monitoring, continuous validation, continuous integration (CI), continuous delivery (CD), continuous deployment, version control, scalability, elasticity

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What type of debugging is carried out by examining the code *without* executing the program? It can be done by scrutinizing code visually, or with the aid of specific automated tools—static code analyzers—based on the language being used.

2. What is the process of measuring your source code's quality when it is passed on to the quality assessment (QA)?
3. What is the software development environment where you create your code on your computer or in the cloud?
4. What is a development environment that allows you to test your code or application but is as similar to the production environment as it can be? This environment allows you to ensure that each component of your application still does its job with everything else going on around it.
5. What is a software development and project management process where a project is managed by breaking it up into several stages and involving constant collaboration with stakeholders and continuous improvement and iteration at every stage?

This page intentionally left blank



This chapter covers the following topics related to Objective 2.4 (Summarize authentication and authorization design concepts) of the CompTIA Security+ SY0-601 certification exam:

- Authentication Methods
  - Directory Services
  - Federation
  - Attestation
  - Technologies
    - Time-based one-time password (TOTP)
    - HMAC-based one-time password (HOTP)
    - Short message service (SMS)
    - Token key
    - Static codes
    - Authentication applications
    - Push notifications
    - Phone call
  - Smart card authentication
- Biometrics
  - Fingerprint
  - Retina
  - Iris
  - Facial
  - Voice
  - Vein
  - Gait analysis
  - False acceptance
  - False rejection
  - Crossover error rate
- Multifactor Authentication (MFA) factors and attributes
  - Factors
    - Something you know
    - Something you have
    - Something you are
  - Attributes
    - Somewhere you are
    - Something you can do
    - Something you exhibit
    - Someone you know
- Authentication, authorization, and accounting (AAA)
- Cloud vs. on-premises requirements

Summarizing Authentication and Authorization Design Concepts

One of the most important design considerations when developing applications is the security model. Adding security controls and capabilities later will cause issues, requiring much more time and effort than if they were added up front. This effort will be a huge undertaking because the security model does not accommodate what the organization required or wanted. In this chapter you learn design concepts and details about authentication and authorization and gain a better understanding of why they need to be built in at the beginning of any development lifecycle.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table12-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 12-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Authentication Methods	1–2
Biometrics	3–4
Multifactor Authentication (MFA) Factors and Attributes	5–6
Authentication, Authorization, and Accounting (AAA)	7–8
Cloud vs. On-premises Requirements	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following best describes a directory structure used on Microsoft Windows-based servers and computers to store data and information about networks and domains?
   a. Token Ring Network
   b. Packet Passing Ring
   c. AAA
   d. Active Directory
   e. TOPD

2. Which of the following should be distributed among different physical locations to provide redundancy for user authentication?
   a. Linux servers
   b. Azure Webservers
   c. Smart cards
   d. Directory services

3. Which of the following is a contactless technology that is capable of matching a human face from a digital image?
   a. Fingerprint authentication
   b. Retinal authentication
   c. Facial recognition
   d. Iris authentication

4. What biometric security measure is characterized by the identification and authentication of vocal modalities?
   a. Retinal authentication
   b. Voice verification
   c. Facial recognition
   d. Fingerprint authentication

5. A physical item like a smart card (CAC) represents what multifactor authentication (MFA) factor or attribute?
    a. Something you know
    b. Something you have
    c. Something you are
    d. Something you exhibit

6. Which of the following is an example of a biometric security measure that is something you exhibit?
    a. Personality trait or behavior
    b. Geolocation or IP address
    c. Facial recognition
    d. Photo ID

7. What does the abbreviation AAA mean?
    a. Authentication, authorization, and accounting
    b. Authorization, accounting, and accredited
    c. Accounting, abbreviate, and accentuate
    d. Accounting, allocation, and attenuate

8. Which of the following is a feature that provides a way of identifying a user using a user ID and a password specific to that user?
    a. The AAA authentication feature
    b. Authorization, accounting, and accredited
    c. Authorization and accounting
    d. Diametric simulation

9. Which method provides highly available, always-on, pay-as-you-grow systems and services?
    a. Enterprise Data Center
    b. HAS highly available services
    c. Cloud-based hosted systems
    d. Next-generation data services

10. From a compliance perspective, which method of authentication hosting is preferred?
    a. Static code placement
    b. Enterprise on-premises hosting
    c. Cloud-based hosted systems
    d. Third-party hosted data and systems

Foundation Topics

Authentication Methods

If you are serious about application security, you must start with a solid understanding of authentication methods. Several components are involved in accomplishing these objectives. Security hinges on two simple goals:

- Keeping unauthorized persons from gaining access to resources
- Ensuring that authorized persons have access to and can access the resources they require

One of the key concepts covered here is authentication and authorization, components that are at the core of the CIA triad, depicted in Figure 12-1. In terms of security, here *CIA* stands for *confidentiality*, *integrity*, *and availability*—three principles that form the cornerstone of any organization's security infrastructure.

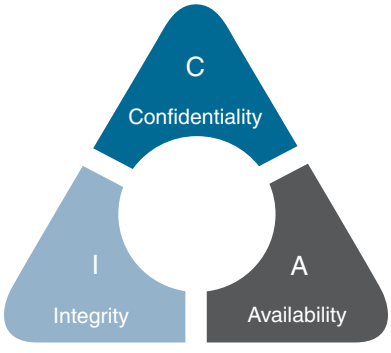


FIGURE 12-1 The CIA Triad

NOTE See the NIST 1800-25 Draft at https://www.nccoe.nist.gov/library/data-integrity-identifying-and-protecting-assets-against-ransomware-and-other-destructive-1 for more information.

TIP Authentication and authorization are often confused with each other. However, although these two elements work closely together, authentication verifies the user's identity, whereas authorization verifies that the user in question has the correct permissions and rights to access the requested resource, as outlined in Figure 12-2.

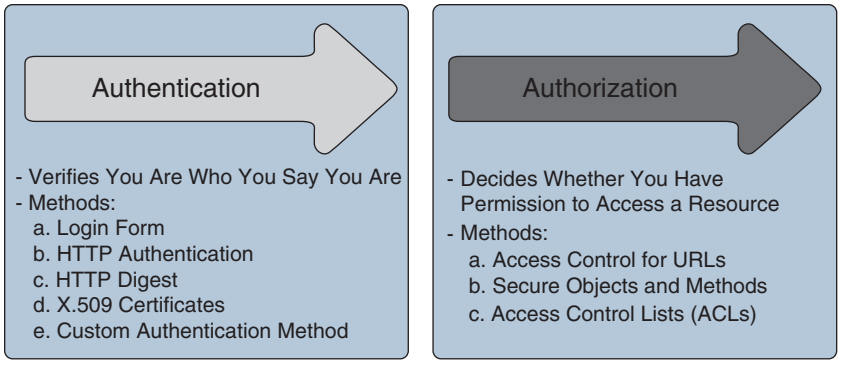


FIGURE 12-2 Authentication vs. Authorization

In the initial step, the authentication system will verify individuals are who they say they are. Provided they are, these individuals receive a pass for future access/admittance. This process is similar to showing a gate agent your ticket at a club, where if you leave, the gate agent stamps your hand or provides a wristband to allow you back in or is used to verify your age. In this way, the stamp is your token. In application security, your ticket is your login credential, and once verified, the token is a unique combination of characters that another program reads. The wristband is symbolic for the "role" granting you specific access privileges.

When a system is verifying an individual, certain actions can be requested on the resource, such as potentially filtering the results. A token proves you are allowed, and the role allows you access to the resource, but the results are a filtered look.

Key Topic

Authentication is the process of an entity (the principle) providing its identity to another entity (the system); it is the first step of a two-step process. There are a variety of ways to authorize users, even entrusting authorization to third parties that are also known as identity providers (IdPs).

Authorization is proving individuals are who they say they are and receiving an endorsement from a trusted party, either your own system or another. Authorization happens repeatedly after someone is permitted to access an application. Every time a page is requested or a button is clicked, an application should check whether that user *can* do what is requested. A remedial approach to authorization is to hard-code role checks within the code like (if %user.role = admin) to display a block of functionality. This approach limits the flexibility of the software in the future, so you must update code everywhere to enable a feature.

Token-based authentication is stateless; no session information is stored on the server. When the user authenticates, the server sends a ***token*** to the client, where it saves the token locally, and when subsequent application access is requested,

the client includes the token in the request and the server validates and sends the response back to the client. The application then continues to work as expected.

Authentication components are considered a preventive measure that can be broken down into three factors and four attributes, which we cover in more detail in "Multi-factor Authentication (MFA) Factors and Attributes."

Key Topic

Directory Services

Directory services are customizable information stores that function as a single point of truth, from which users can locate resources and services throughout the network. This customizable information store also gives you, as administrator, a single point for managing objects and their attributes. Many types of directory services are distributed to make the service highly available.

The most common directory services are Microsoft Active Directory (AD) and LDAP. For the purpose of this discussion, Azure Active Directory has the same features as Microsoft Active Directory. The authentication process starts with a user communicating with Active Directory #1. Provided the user is authentic, Active Directory provides the authenticated user a token pair #2, and finally the user can access the service or resource he or she has authorization for #3. As illustrated in Figure 12-3, Active Directory does not usually implement the transaction or rollback schemes that regular databases require. If they are permitted at all, directory updates are typically simple all-or-nothing changes. Directories are tuned to respond quickly to high-volume lookup or search operations. A lookup is an operation that targets a specific, unique entry, such as a domain name. A search is an operation that targets data common to multiple entries, such as the data collected by an Internet search engine about a topic.

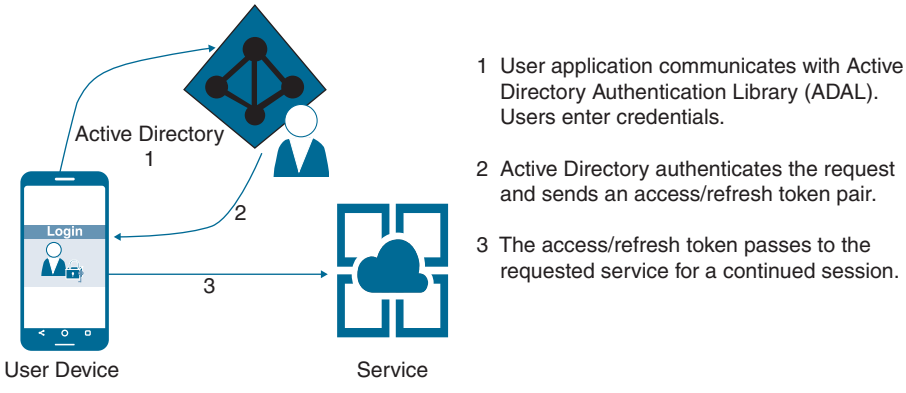


FIGURE 12-3 Active Directory Authentication Process

Attacks on directory services have long been documented, along with remediation and mitigation techniques. Legacy protocol attacks such as NetBIOS NS, LLMNR, and WPAD can easily be disabled with GPO policy actions such as making a change to the registry under \Computer Configuration\Administrative Templates\Network\DNS Client\Turn off Multicast Name Resolution.

Other attacks such as weak passwords, excessive privileges, and unencrypted credentials in memory can easily be mitigated with the appropriate enforcement and settings. Kerberos attacks such as Kerberoasting TGS can be mitigated with strong service account passwords, disabling RC4 encryption, and removing unused service accounts.

Federations

Federations are considered identity providers (IdPs). A federation is a process wherein one system is responsible for authentication of a user, and that systems then sends a message to a second system announcing who the user is and verifies that the user was properly authenticated.

The most widely used federation mechanism today is Security Assertion Markup Language (SAML). SAML is an older open standard that allows IdPs to pass authorization credentials to service providers (SPs). OpenID Connect, which is a simple identity authentication protocol built on top of OAuth 2.0, is becoming more common as a replacement for SAML. OAuth provides resource authorization and was created by Google, Twitter, and others to address SAML's lack of mobile platform capabilities.

Federation allows single sign-on (SSO) without passwords. SSO is a technology that combines several different application logins into one. With SSO, users have to enter their login credentials (username and password) only one time on a single page to access all applications. In Figure 12-4, the federation server knows the username for a person in each application and presents that application with a token that says, "This person is domain\janesmith or janesmith@example.com." No password is required for the user to log in to each system. Because of the trust between the two systems, the target application accepts this token and authenticates the user. You have likely seen this in Facebook, Gmail, Twitter, LinkedIn, and others, where it prompts you to use one of these methods to authenticate to another service.

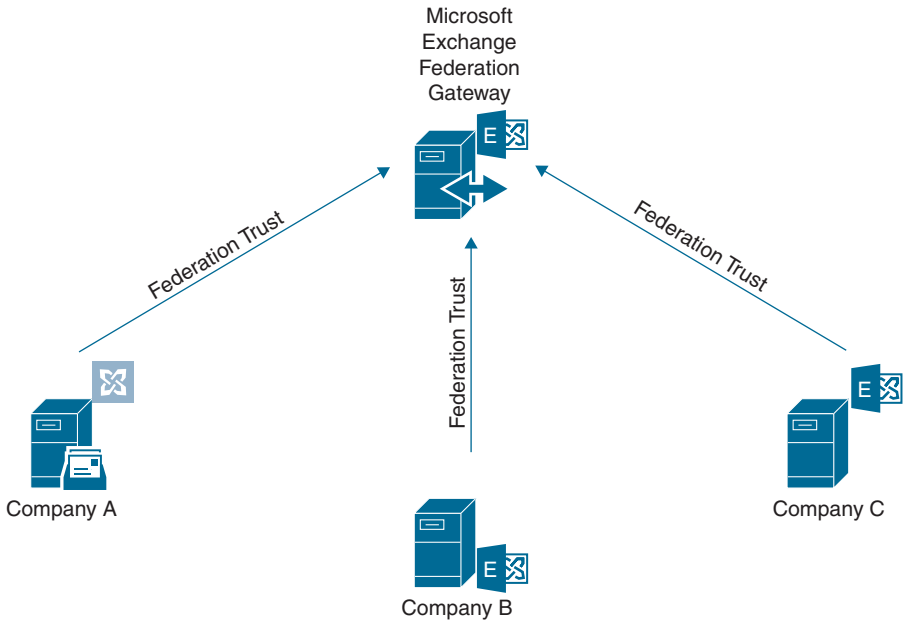


FIGURE 12-4 Federation Interactions

Attacking IdPs start by exploiting SSO. IdPs are considered trusted third parties, so when attackers are allowed to configure their own IdP and connect to unsuspecting targets, they could compromise the security of all accounts on the SP for which they were configured. There are two known attack vectors: ID spoofing and key confusion. The real problem here is the trust assumptions of different components of the system. To mitigate this attack from locking down systems, you should ensure that you are working only with trusted IdPs, enable consistent updating and patching of the systems you are responsible for to have the latest libraries and binaries, and ensure your logging and monitoring are continuously reviewed.

Attacking the federation—and we are not speaking about *Star Trek* here—usually starts with exploiting Active Directory, which we covered earlier. Another method is a Golden SAML attack, which starts with a forged SAML token. It bypasses MFA requirements imposed by applications. Performing any number of the listed attacks would allow an attacker to impersonate an already authenticated user, and using preauthentication, would proceed to company B and C unopposed, with rights and privileges associated with that account on those systems and those granted by the federations.

Attestation

Attestation provides evidence or proof of something. The process of attestation allows one program to authenticate itself, where remote attestation is a means for one system to make reliable statements about the software it is running to another system. The remote party can then make authorization decisions based on that information.

NOTE To find out more about attestation, see www.w3.org/TR/webauthn/.

Attestation statement formats are identified by a string, called an attestation statement format identifier, chosen by the author of the attestation statement format.

Attestation statement format identifiers *should* be registered per WebAuthn-Registries, "Registries for Web Authentication". All registered attestation statement format identifiers are unique among themselves as a matter of course.

When the key pair is created, there is an option to request the attestation certificate. This optional information can be sent to the relying party as part of the registration process. The attestation is how authenticators prove to the relying party that the keys they generate originate from a genuine device with certified characteristics and establish a hardware root of trust.

The purpose of attestation is to cryptographically prove that a newly generated key pair came from a specific device. This process provides a root of trust for a newly generated key pair and enables you to identify the attributes of a device being used, how the private key is protected, and what kind of biometric or Trusted Platform Module (TPM) authentication system is being used.

A TPM Quote operation is used to authoritatively verify the contents of a TPM's platform configuration registers (PCRs). During provisioning, a composite hash of a selected set of PCRs is computed. The TPM Quote operation produces a composite hash that can be compared with the one computed while provisioning.

Certain implementations of attestation are susceptible to replay, masquerading, tampering, and hardware attacks. A malicious attesting system can replay old values of measurements that correspond to a valid response before the attesting system became corrupted. In a masquerade attack, the attacker can send measurement lists and a TPM quote of another valid system to get the attacked system to trust it. To protect against these types of attacks, you should make sure you connect only to trusted third parties in which you have audit, logging, or certification of vulnerability tests performed regularly. You also should configure higher levels of logging on your own systems to catch abuse.

Authentication Methods and Technologies

Cybercriminals are always improving their attack methods and processes. As a result, security teams are facing plenty of authentication-related challenges. This is why companies are starting to implement more sophisticated authentication technologies as part of the process. The following sections describe some common authentication methods used to secure modern applications and systems, including one-time passwords, multifactor authentication, tokens, and phone-based systems.

Time-Based One-Time Password (TOTP)

A ***time-based one-time password (TOTP)*** is a one-time password (OTP), a temporary passcode generated by an algorithm that uses the current time of day as one of its authentication factors (hence the temporariness). These passwords are commonly used for two-factor authentication (2FA) and have been adopted by many cloud applications and services. In this context, a user utilizes a two-factor authentication method; the user must enter a traditional potentially static password as well as a TOTP to gain access to email, a process similar to Gmail two-factor authentication. Typically, the temporary passcode expires after a certain amount of time, usually after a few minutes.

In two-factor authentication, for example, the user logs in to a bank account with a username and password, and an SMS message or email with the time-based code is sent to the user for input into the banking access form, prior to allowing the user to log in. The username and password are known to the user, and the TOTP is sent to a device the user owns or has access to.

TOTP attacks usually begin with a phishing attack, where the attacker impersonates a certain service or website and asks the user for the TOTP code, which allows the attacker to grab the code and immediately use it to log in as the victim. MITM attacks also occur if an attacker is able to sniff unencrypted traffic. If the TOTP device is a phone or token and is stolen or out of the owner's control for a short time, this can also lead to a breach.

HMAC-Based One-Time Password (HOTP)

The ***HMAC-based one-time password (HOTP)*** algorithm is based on hash-based message authentication codes (HMAC). It is the foundation for open authentication (OAuth). HOTP was published as an IETF RFC 4226 in December 2005.

The HOTP algorithm provides a method of authentication by symmetric generation of human-readable passwords, or values, each used for *only one authentication attempt*. The one-time property is directly from the single use of each counter value.

Parameters must be established if intending to use HOTP; typically, these are specified by the authenticator, and either accepted or not by the authenticated: both parties compute the HOTP value derived from the secret key K and the counter C. Then the authenticator checks its locally generated value against the value supplied by the authenticated.

The authenticator and the authenticated increment the counter, C, independently of each other, where the latter may increase ahead of the former; thus, a resynchronization protocol is wise. This simply has the authenticator repeatedly try verification ahead of their counter through a communication window size. The authenticator's counter continues forward of the value at which verification succeeds and requires no actions by the authenticated.

There is a recommendation that persistent throttling of HOTP value verification take place, to address their relatively small size and thus vulnerability to brute-force attacks. It is suggested that verification be locked out after a small number of failed attempts or that each failed attempt attracts an additional linearly increasing delay.

After verification, the authenticator can authenticate itself simply by generating the next HOTP value, returning it, and then the authenticated can generate its own HOTP value to verify it. Note that counters are guaranteed to be synchronized at this point in the process.

HOTP is suspectable to brute-force attacks, keyboard logging, on-path attacks, and phishing. HTOP codes are valid until they are used. The attacker does not need access to your OTP token device, only the area or system where you may have stored a code or utilized a computer-based OTP. Since the hash value is a predetermined size, it is mostly resistant to calculation attacks.

NOTE For more information, see www.ietf.org/rfc/rfc4226.txt.

Short Message Service (SMS)

Short Message Service (SMS) is a method used to provide one-time passwords to different types of OTP services.

Mobile devices with SMS text capability can be used for authentication via OTP and challenge/response, where you are asked to answer Yes or No (CR or Y/N). This is a less-secure form of strong authentication because it is vulnerable to on-path attacks (previously known as man-in-the-middle, or MITM, attacks). SMS OTP sends an OTP to the user's phone via SMS, and the user enters the OTP into the login authentication and is approved.

SMS challenge/response sends a question asking if the authorization attempt is approved to the user's phone via SMS. If the user texts back "Yes," authentication is completed, and the user is logged in. If the user texts back "No," authentication fails, and the user is not logged in. SMS is suspectable to SIM takeover, brute-force attacks, keyboard logging, on-path attacks, phishing, and common computer access attacks because, after all, most cell phones are computers.

Token Key

A ***token key*** can be a hardware or software device that on regular intervals changes the code displayed or provides a new key code each time pressed and is based on a certain algorithm. One of the most common types that people remember is the RSA token key (see Figure 12-5).



FIGURE 12-5 RSA Token Key

Token-based authentication is a security technique that uses a hardware security token to provide a code to the user that expires after 60 seconds. The user must enter the code within the allotted time or has to start over. Most of the older-style token keys relied on a vulnerable algorithm that was defeated, which made them obsolete. A newer method is Universal 2nd Factors (U2F), which uses a new standard in authentication called Fast Identity Online (FIDO).

Newer versions of hardware authenticator solutions support FIDO2, which replaces weak password-based authentication. One of the more common tokens used today is the YubiKey strong two-factor authentication USB token key. Some of the newer ones require a fingerprint to activate.

NOTE To find out more about YubiKey and strong two-factor authentication devices, see www.yubico.com/products/.

The challenges with hardware tokens are that they are expensive, easy to lose, and vulnerable to theft and breach of codes; they also require constant administration and maintenance.

Static Codes

When we talk about standard passwords, we think of static codes; these are security codes you change based on predetermined policies. Static codes are among the weakest and most insecure methods for authenticating. They should be used only as a backup mechanism and are never considered secure because they are available in some physical form, whether on paper or stored digitally. Many consider them a last resort method to gain access (in the case of a lost password). However, static codes can be used with other authentication types such as two-factor authentication (2FA), where the user enters a username and password and then is prompted for a one-time use static code. Static codes are easily hacked, and no matter how strong you make them, there are dictionaries and techniques for cracking even the strongest passwords.

Authentication Applications

Authentication applications consist of an authenticator and a verifier. A number of authentication mechanisms can be used:

- Basic, single-factor authentication, usually built into the application, is a single credential pair, such as a user ID and password.
- Multifactor authentication is a combination of authentication methods to validate the identity; usually starts with a user ID and password and is supplemented with a smart card, a token, or pin.
- Cryptographic authentication involves the use of cryptography, in the form of a public key, digital signatures, certificates, and authentication codes.

Perhaps you are currently using Gmail for personal email. In that case, you might have even enabled the authenticator (Gmail calls this process two-step verification) when attempting to log in to your email and enter your username and password combination. You then use the Google authenticator and obtain a code that is good for 60 seconds—enough time for you to enter and gain access to your account. Microsoft email has a similar authenticator that operates in the same manner.

NOTE You can find more information about Google two-step verification at www.google.com/landing/2step/.

Other popular two-step verification mechanisms and software are Cisco Duo, Auth0, Okta, and Jamf. Many others have multiple features, including MFA, 2FA, SMS, and Push.

Push Notifications

Push notifications are brief messages or alerts that are sent through an application or the telco SMS system discussed earlier. A user receives an authentication verification prompt and then is prompted to approve or decline access to a certain resource or to log in to a system. Notifications can be sent in-band or out-of-band using any number of communication channels. Push notifications authenticate the user by confirming that the device registered with the authentication system, typically a mobile device, is in fact in the user's possession at the time of the request. When you register your account, you link a mobile device you own, with your username or ID. Some implementations do not require a password at all and simply prompt the user on the device to approve. Google's Gmail and Microsoft Outlook email have the capability to send push notifications. One concern is if someone clones your phone's SIM chip or if an attacker is eavesdropping by monitoring the telco network. In that case, the attacker can intercept these types of authentication messages and gain access to your account. We have seen these sorts of attacks play out successfully where attackers have stolen millions in Bitcoin. These solutions are implemented as vendor specific for their own services, which makes them unavailable for an organization- or enterprise-wide solution. Also, most of these solutions are used or offered as secondary authentication methods only. One such push type notification application is DUO. It can be configured to send you a yes or no message to allow a user to perform two-step (MFA) authentication using your cell phone.

Phone Call Authentication

Phone call authentication takes place where an automated system calls a predefined phone number that was registered when set up. The automated system can read off a number/letter combination that the user enters into the authenticator, thereby providing another means of identifying the user. This is considered a multifactor authentication method and "something you have," such as the authenticator. Other methods include a phone call prompting the user to enter a certain number combination into the phone to be compared and approved, as well as prompting the user to press a single number or hash/star to confirm he or she requested the call. Some password reset systems such as Microsoft's implementation require three levels of verification:

1. You request the reset.
2. You verify the email.
3. An automated call requests you enter the code that was emailed or texted to you to reset the password.

Another method is displaying an image of something you previously selected. If you enter it wrong, the system will continue to lock you out of the account.
The methods for defending against these types of attacks are similar to those of

SMS-based authentication: You first secure your SIM card. Most telcos allow you to log in to your account and lock the SIM within their service, as well as lock the SIM on the phone itself. You also should keep your phone up to date on security patches and updates and run phone-specific antivirus software.

Smart Card Authentication

A ***smart card*** is a secure microcontroller that is typically used for generating, storing, and operating on cryptographic keys. Smart card authentication provides users with smart card devices for the purpose of authentication. Users connect their smart card to a host computer. Software on the host computer interacts with the key's material and other secrets stored on the smart card to authenticate the user. For the smart card to operate, a user needs to unlock it with a user PIN. Smart cards now come with RFID/NFC capabilities to allow users to simply tap the device they want to authenticate against, such as a cell phone or computer. The software on the device then is activated and reviews the data and either allows or declines the user to authenticate based on data on the card matching data stored on the device or a remote or cloud application.

Key Topic Biometrics

Biometric identifiers are distinctive, measurable body measurements and calculations related to human characteristics. Biometrics are used for authentication and a form of identification and access control. Many different aspects of human physiology, chemistry, or behavior can be used for biometric authentication, as detailed in the sections that follow. Common biometrics are fingerprint readers that use mathematical representations of your fingerprint to authenticate you. Iris, retinal, and eye authentication requires you to look into a device that reads certain parts of your eye, whereas facial recognition scans your face and identifies certain points of your face. There is also voice recognition, like in the old movie *Sneakers* ("My voice is my passport. Verify me."), where your voiceprint is used and compared to a similar mathematical or audio analysis format that authenticates you. As with many other authentication methods, attacks on biometric systems have been successfully executed over the years. Error rates allow attackers to come up with creative new tactics to bypass these methods.

Fingerprints

One of the oldest and best known biometric methods for authentication is fingerprints. Fingerprint scanners have two jobs: first, to obtain an image of your finger, and second, to determine whether the patterns of ridges and valleys in the image match the patterns of prescanned images. Only specific characteristics that are

unique to every fingerprint are filtered and saved as an encrypted biometric key, usually in a mathematical representation. No fingerprint is ever saved—only a series of binary codes. Fingerprints are easy to bypass and circumvent when attackers have physical access to any of the user's items. The mathematical representation is typically stored, so if an attacker gains access to that data, they do not have the actual fingerprint, only a digital representation. This occurred in the U.S. Office of Personnel Management (OPM) attack in 2013, when unknown hackers infiltrated the systems of a government contractor and stole personnel files that included credentials of all government employees.

TIP You can read and refer to the latest NIST ANSI/NIST-ITL projects on fingerprints, along with other methods and standards, at www.nist.gov/programs-projects/fingerprint.

Retina

Scanning the retina of the eye is one of the most well-known and the least deployed biometric authentication methods, mainly due to costs. A retinal scanner scans the unique patterns of a person's retina. The blood vessels within the retina absorb light more readily than the surrounding tissue and are easily identified with proper lighting. The retinal scan is performed by casting an unperceived beam of low-energy infrared light into the person's eye as he or she looks through the scanner eyepiece.

TIP You can read about and refer to the latest on retina benchmarks from the American Academy of Ophthalmology at https://www.aao.org/summary-benchmark-detail/retina-summary-benchmarks-2020.

Iris

Iris recognition is an automated method of biometric identification that uses mathematical pattern-recognition techniques with images of one or both of the irises of an individual's eyes, whose complex patterns are unique, stable, and can be seen under normal lighting conditions.

Facial

Facial recognition systems are a contactless technology that is capable of matching a human face from a digital image or a video frame against a database of faces. The most advanced face recognition method, which is also employed to authenticate users through ID verification services, works by pinpointing and measuring facial features from a given image.

Voice

Voice and speech recognition are two separate biometric modalities that, because they are dependent on the human voice, can vary or be inconsistent. Both are contactless, software-based technologies, and as such are counted among the most convenient biometrics in regular use. Voice recognition, also commonly referred to as a voiceprint, is the identification and authentication of vocal modalities. By measuring the sounds a user makes while speaking, voice recognition software can measure the unique biological factors that, combined, produce his or her voice.

Vein

Vein authentication can take the form of the veins in the person's hand or palm, using blood vessel patterns as a personal identifying factor. The vein information is hard to duplicate because a palm has a broader and more complicated vascular pattern.

Gait Analysis

Gait analysis enables a system to authenticate a person from a long distance using low-resolution machine vision-based systems. The information and properties regarding a person's "gait" are extracted and compared with stored samples, which in turn make authentication possible. Essentially, this is a behaviorally biometric modality that identifies a human on the idiosyncratic way of walking. This is the study of a human motion, body mechanics, and the activity of the muscles, measuring body movements and body mechanics. While difficult, these systems are not foolproof. Depending on the height, stride, and pace, it is possible through practice to mimic a person's gait. These systems have been mainly used as a tertiary method of verification more than authentication.

Efficacy Rates

Efficacy rates vary depending on the technology. As newer and more accurate methods for using biometrics as an identity provider expand, so do the efficacy rates. For example, certain eye diseases and metabolic conditions can reduce or negate the efficacy of an eye or iris scan. Veins might not be as pronounced in certain people with different types of skin or skin diseases. Social and cultural considerations must also be taken into consideration because some people might not like touching devices in fear of contracting diseases.

False Acceptance

The ***false acceptance rate (FAR)*** is the measurement of the likelihood that a biometric security system will incorrectly accept an access attempt by an unauthorized user. A system's FAR typically is stated as the ratio of the number of false acceptances divided by the number of identification attempts. False rejection rate (FRR) is the percentage of identification instances in which false rejection occurs. The crossover error rate (CER) is the percentage at which the FAR and FRR are equal.

False Rejection

The ***false rejection rate (FRR)*** is the measure of the likelihood that the biometric security system will incorrectly reject an access attempt by an authorized user. Figure 12-6 shows that as the number of false acceptances goes down, the number of false rejections will go up and vice versa. The point at which the lines intersect is the equal error rate (EER), sometimes called the equilibrium, which is where the percentage of false acceptances and false rejections is the same. If you try to reduce the FAR to the lowest possible level, the FRR will likely rise sharply. The more secure your access control, the less convenient it will be, as users are falsely rejected by the system. The same also applies the other way round. To increase user convenience, you reduce the FRR, so in turn the system is likely to be less secure with a higher FAR.

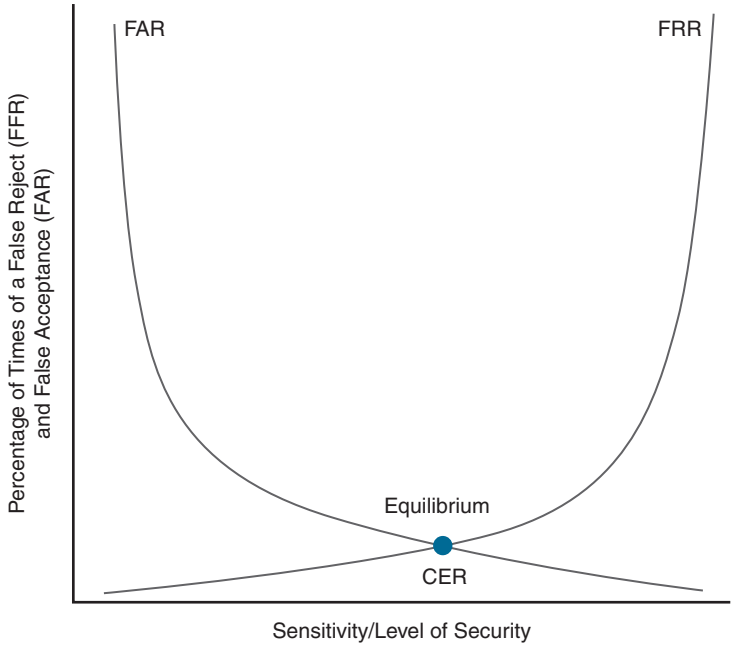


FIGURE 12-6 Equal Error Rate (EER)

Crossover Error Rate

The Crossover error rate describes the point where the False reject rate (FRR) and false accept rate (FAR) are equal. CER is also known as the equal error rate (EER). Authentication algorithms need to simultaneously minimize permeability to intruders and maximize the comfort level; therefore, they must be both demanding and permissive. This contradiction is the base for the optimization problem in authentication algorithms, and the measure of success for the overall precision of an algorithm and of its usability is the ***crossover error rate (CER)***, the error rate obtained at the threshold that provides the same false acceptance rate and false rejection rate. Simply put, the CER is the percentage at which the FAR and FRR are equal.

Key Topic Multifactor Authentication (MFA) Factors and Attributes

Multifactor authentication (MFA) is the process of identifying a user by validating two or more claims presented by the user, each from a different category of factors. It is sometimes called two-factor authentication (2FA), in which a user is required to present more than one type of evidence in order to authenticate on a system. When we talk about multifactor authentication, the major factors are based on (1) your memory, (2) your possessions, and (3) yourself. Take a moment and review each, and reflect on how your personal traits apply to each of them.

- **MFA Factors:**
  - **Something you know:** A user ID, password, passphrase, PIN code, secret code, or security questions.
  - **Something you have:** A physical item like a smart card (CAC), RFID card, or token, certificate, email, SMS, or phone call.
  - **Something you are:** Biometrics such as your handprint, fingerprint, eyes or retina, facial recognition, and voiceprint.

Authenticating with a user ID and a password is considered single-factor authentication because it is a single-factor method. MFA factors are really authentication factors; an authentication factor is a category of credential that is used for identification purposes and identity verification. For MFA, each additional factor is intended to increase assurance that a "person" or "system" involved in that request or communication that is making the request to access a system is who or what they are stated to be. *Something you know* is considered a knowledge factor; in the context of security, it's something the user possesses such as a personal identification number (PIN), a username, a password, or an answer to a secret question. You have likely been asked to set up four or five questions to access your bank account, and when logging in, you are prompted to answer one of them. *Something you have* is typically a physical object in your possession, such as a smart card, security token, USB token, even your

ATM card, or key. Figure 12-7 shows multifactor authentication—something you are, you know, and you have—and where they intersect to create dual- and three-factor authentications.

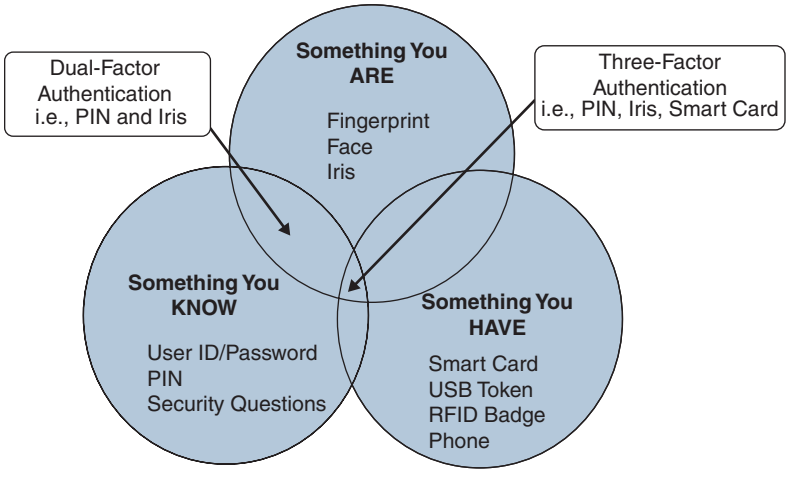


FIGURE 12-7 Multifactor Authentication

For example a user goes to the bank and uses the automated teller machine (ATM), enters her ATM card into the machine, and then is prompted for a PIN. After she enters the PIN correctly, she has full access to her account to withdraw money (only a certain dollar amount). This is an example of dual-factor authentication, requiring both something you have (physical key; ATM card) and something you know (PIN). If this same person goes into the bank and wants to withdraw more than the ATM allows, she must provide her withdrawal slip, an ID where the teller verifies it looks like her, and an account number from the paper, or she must insert her card in the indoor terminal, which prompts her to enter her PIN. This scenario is still considered a dual-factor authentication method because it is all part of the same method; the teller is not making a machine-scan verification match of the person's face.

An attribute is something such as a physical characteristic, something you can do, such as your signature, the way you type, mistakes, patterns, also the way you walk, where you're located, and through a third party vouching that they know you are who you claim to be. Review the following multifactor authentication attributes:

- **MFA Attributes:**
  - **Something you can do:** Accurately reproducing a signature, writing technique, typing technique
  - **Something you exhibit:** Personality trait, behavior

- **Someone you know:** Authentication by chain of trust
- **Somewhere you are:** Geolocation, IP address, where you're located

Key Topic Authentication, Authorization, and Accounting (AAA)

Authentication, authorization, and accounting (AAA) is the framework for intelligently controlling access to computer resources, enforcing policies, auditing usage, and providing the information necessary for audits.

This process starts with identification, proving you are who you are. *Authentication* processes provide a way of identifying a user, typically by having the user enter a valid username and valid password before access is granted. The AAA server compares a user's authentication credentials with other user credentials stored in a database, which could be Active Directory. If the credentials match, the user is granted access to the network. If they do not match, the authentication fails, and access to the network is denied. Password reuse has been one of attackers' favorite methods to gain access to systems controlled by AAA. If a user's credentials are exposed on one platform or even made public, attackers can add these passwords and usernames to their attack dictionary.

Authentication supposes a user wants to gain *authorization* to a certain resource. After logging in to a system, for instance, the user may try to access a network drive share, printer, or gateway. The authorization process determines whether the user has the authority to issue such commands. In other words, authorization is the process of enforcing policies—determining what types or qualities of activities, resources, or services a user is permitted. Usually, authorization occurs within the context of authentication. After you have authenticated a user, that user may be authorized for different types of access or activity.

The final *A* of the AAA framework is *accounting*, which logs the resources a user consumes during access. This may include the amount of system time or the amount of data a user has sent or received during a session. Accounting is carried out by logging session statistics and usage information and is used for authorization control, compliance, billing, trend analysis, resource utilization, and capacity planning activities.

Key Topic Cloud vs. On-premises Requirements

Cloud-based authentication typically employs third-party managed platforms and usually includes application programming interfaces (APIs) that use the OAuth 2.0 protocol for authenticating both user accounts and service accounts. The benefit of these third-party managed cloud authentication services is availability, flexibility,

and speed. These systems are generally highly available and allow users all over the world to authenticate. Many companies are already using cloud authentication services, such as when you access a SaaS or PaaS application or service. Although they can be tied to local directories, it is more cost effective to utilize authentication services entirely from the cloud provider; however, it isn't as scalable.

With *on-premises* authentication systems, a company hosts its own system in its data center, along with all the expenses associated with the compute system, including power, space, and cooling. These systems are available only as long as the company's network remains available. Any external users requiring access and authentication will require additional administration; that includes granting and managing users and access by local staff.

Table 12-2 provides a comparison of cloud versus on-premises requirements.

Table 12-2 Cloud vs. On-premises Requirements

Capability	Cloud	On-Premises
Deployment	Resources are hosted in one or more service provider environments, where the systems are managed by the service provider.	Resources are deployed in-house with an enterprise's existing IT infrastructure. The company must have expertise on staff to monitor and manage the local systems.
Cost	Enterprises pay only for what they use, with low or no maintenance or upkeep costs.	Enterprises are responsible for the ongoing costs, server hardware, power, space, and cooling environments.
Control	Data ownership or disposition is questionable. Unexpected issues and downtime could cause access issues.	Enterprises retain full control of their data and what happens to it.
Security	Misconfigured resources have often been the biggest concern, and have been improving with more granular controls.	Enterprises are extra sensitive to this information and put extra controls and monitoring in place to protect assets.
Compliance	Enterprises are required to ensure cloud-based authorization continuously meets regulatory requirements and provide proof that sensitive data is secured.	Enterprises that are required to meet compliance requirements may choose to host everything in-house to ensure they remain compliant.

TIP Unlike on-premises solutions, cloud-based solutions are mostly subscription based.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 12-3 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 12-3 Key Topics for Chapter 12

Key Topic Element	Description	Page Number
Paragraph	Defining authentication	290
Section	Directory Services	291
Section	Biometrics	300
Section	Multifactor Authentication (MFA) Factors and Attributes	304
Section	Authentication, Authorization, and Accounting (AAA)	306
Section	Cloud vs. On-Premises Requirements	306

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

authentication, authorization, token, directory services, attestation, time-based one-time password (TOTP), HMAC-based one-time password (HOTP), Short Message Service (SMS), token key, authentication applications, push notifications, phone call authentication, smart card, Biometrics, false acceptance rate (FAR), false rejection rate (FRR), crossover error rate (CER)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What can replicate data widely to increase availability and reliability and thus reduce response time?

2. What directory service operation targets a specific, unique entry, such as a domain name?
3. What is the process where one system is responsible for authentication of a user and provides that information to another resource as authenticated?
4. Which biometric method uses blood vessel patterns as a personal identifying factor?
5. What is the study of a human motion, body mechanics, and activity of the muscles?
6. What is the point where the false rejection rate (FRR) and the false acceptance rate (FAR) are equal?



This chapter covers the following topics related to Objective 2.5 (Given a scenario, implement cybersecurity resilience) of the CompTIA Security+ SY0-601 certification exam:

- Redundancy
  - Geographic dispersal
  - Disk
    - Redundant array of inexpensive disks (RAID) levels
    - Multipath
  - Network
    - Load balancers
    - Network interface card (NIC) teaming
  - Power
    - Uninterruptible power supply (UPS)
    - Generator
    - Dual supply
    - Managed power distribution units (PDUs)
- Replication
  - Storage area network
  - VM
- On-premises vs. cloud
- Backup types
  - Full
  - Incremental
  - Snapshot
  - Differential
  - Tape
  - Disk
  - Copy
  - Network-attached storage (NAS)
  - Storage area network (SAN)
  - Cloud
  - Image
  - Online vs. offline
  - Offsite storage
    - Distance considerations
- Non-persistence
  - Revert to known state
  - Last known-good configuration
  - Live boot media
- High availability
  - Scalability
- Restoration order
- Diversity
  - Technologies
  - Vendors
  - Crypto
  - Controls

Implementing Cybersecurity Resilience

Cybersecurity resilience is the capability of an organization to prepare, respond, and recover when cyber attacks happen. An organization has cybersecurity resilience when it can defend itself against these attacks, limit the effects of a security incident, and guarantee the continuity of its operation during and after the attacks. Organizations today are beginning to complement their cybersecurity strategies with cyber resilience. Although the main aim of cybersecurity is to protect information technology and systems, cyber resilience focuses more on making sure the business is delivered.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 13-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 13-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Redundancy	1–5
Replication	6
On-premises vs. Cloud	7
Backup Types	8
Non-persistence	9
High Availability	10
Restoration Order	11
Diversity	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is one of the primary goals of geographic dispersal of data and data processing?
   a. Ensuring availability of data in the event of a disaster or regional issue, such as power outage
   b. Reducing the cost of acquiring computer systems, servers, network infrastructure, and software used for information security
   c. Ensuring the application infrastructure does not affect the underlying network
   d. All of these answers are correct.

2. Which of the following describes the disk redundancy method of RAID 5?
   a. It uses more hard disks and stripes data and parity over all hard disks.
   b. It is a parity hard disk proportional to the log of HDD numbers.
   c. It is a JBOD, otherwise known as just a bunch of drives.
   d. It duplicates data from one hard disk to another; typically, it is two hard disks set up to copy data to both disks.

3. Which hard disk resilience mechanism uses two hard disks to provide redundancy and stripes everything to both hard disks?
   a. RAID 0
   b. RAID 3
   c. RAID 6
   d. RAID 10

4. When a network interface card fails in a server, a secondary NIC continues transmitting traffic uninterrupted. What mechanism allows this feature to operate?
   a. Redundant NIC
   b. NIC teaming
   c. NIC failover
   d. Mirrored NIC

5. Which high-availability power mechanism provides an additional power supply that can share or split the power load?
   a. Redundant supervisor modules
   b. Redundant power supplies
   c. Dual hard disks
   d. RAID 5 power sources

6. SAN-connected servers utilize a special fiber interface card called what?
   a. Network interface card (NIC)
   b. Host bus adapter (HBA)
   c. Storage area network device (SAND)
   d. Data beam fiber (DBF) card

7. You can host services in your own data center or in the cloud. What is considered a cloud-based service?
   a. Azure Active Directory (AD)
   b. Novell NDS Services
   c. Outlook 2010
   d. Data Center Hosted SAP

8. Which backup method provides connected backup storage anywhere in the world with private/public remote storage?
   a. Incremental offline
   b. Intermittent even days
   c. Local backup session
   d. Cloud-based backup
   e. Tape backup to mainframe

9. You can now back up your organization's critical data to the cloud, including all your desktop computers and servers. What is one downside to hosting all your backups in the cloud?
   a. Restoral is available only during the work week.
   b. They require a special VPN connection to each device.
   c. Accessing them may be slow depending on your Internet connection.
   d. Backups are rotated every 30 days.

10. What is the uptime measurement of available time for a highly available system, knowing that no system can be up forever?
    a. 95.99 percent
    b. 100 percent
    c. 90 percent
    d. 99.999 percent

11. Your company just had a catastrophic failure, and the entire data center and campus are down. After you establish your chain of command, what is the first service or system you need to bring up in order to start recovery?
    a. Video telecom VTC systems and bridges
    b. Desktop computers in network operations center
    c. Active Directory and LDAP to authorized users
    d. Network connectivity

12. An organization can add many levels of diversity to enhance its systems and network resilience. Which components will help? (Select all that apply.)
    a. Redundant Pathed Facility Power
    b. Network interface card (NIC) teaming
    c. Supplier and supply chain contracts
    d. Security cameras at all junctions

Foundation Topics

Redundancy

When systems are not designed or configured with redundancy in mind, a single point of failure can bring an entire company to its knees. High availability is provided through *redundant* systems and fault tolerance. If the primary method used to store data, transfer information, or perform other operations fails, then a secondary method is used to continue providing services. Redundancy ensures that systems are always available in one way or another, and with minimal downtime. Several key elements are used in the redundancy process:

- ***Geographical dispersal*:** Placing valuable data assets around the city, state, country, or world can provide an extra level of protection.
- ***Disk redundancy*:** Various levels of disk redundancy exist to reduce data loss from disk failure.
- ***Network redundancy*:** Networks deliver data to users during component failures.
- **Power:** Power resilience provides power via multiple sources.

Geographic Dispersal

Key Topic

We live in a connected world, with employees, contractors, servers, data, and connected resources dispersed throughout the world. This dispersal is essential for a company with a global presence (marketing); it also lends itself to resilience. Local issues are likely to be isolated, allowing the rest of the workforce to continue operating as normal. Companies that have workforces dispersed in remote locations throughout the world have adopted and learned to utilize technology to connect people and resources. To implement geographical dispersal within your company for compute resources, you can set up a computing environment, such as a data center or resource center, in another location at least fifty miles away from the main compute campus, data center. You also should consider placing resources closer to groups of users who might be in another city or country. By way of a wide-area network, you can configure resources that are local to the users to respond first, thereby providing resilience and resource dispersal.

Disk Redundancy

Disk redundancy is simply writing data to two or more disks at the same time. Having the same data stored on separate disks enables you to recover the data in the

event of a disk failure without resorting to expensive or time-consuming data recovery techniques. Implementing disk redundancy starts with selecting the appropriate level of RAID and a controller that is capable of performing RAID functions at various levels. Newer servers have RAID-capable controllers, thus allowing you to develop a RAID strategy and implement the appropriate level of RAID.

Key Topic

Redundant Array of Inexpensive Disks

Redundant Array of Inexpensive Disks (RAID) 1 and RAID 5 are the most common approaches to disk redundancy. RAID is a group of disks or solid-state drives (SSDs) that increases performance or provides fault tolerance or both. RAID uses two or more physical drives and a RAID controller, which is plugged into a motherboard that does not have RAID circuits. Today, most motherboards have built-in RAID, but not necessarily every RAID level of configuration is available. In the past, RAID was also accomplished through software only, but was much slower. Table 13-2 lists a few of the most common levels of RAID and their safeguards, configurations, and protections. Implementing the appropriate RAID level begins with an overall corporate strategy.

Table 13-2 Common RAID Levels and Characteristics

Raid Level	Description	Data Reliability/Speed	Protection
0—Striping	Disk striping; divides data across a set of hard disks.	Is lower than a single disk; enables writes and reads to be done more quickly.	Very low; if a disk fails, you could lose all your data.
1—Mirroring	Duplicates data from one hard disk to another, typically two hard disks.	Provides redundancy of all data because it is duplicated; is much slower.	A complete copy of data is on both hard disks.
5—Striping with parity	Uses more hard disks and stripes data and parity over all hard disks.	Provides higher performance; quickly writes data to all disks.	Data is protected by the parity writes.
6—Striping with dual parity	Similar to RAID 5; includes a second parity distributed across hard disks.	Provides slower storage system performance; provides best protection.	Dual parity protects against loss of two hard disks.
10—Mirroring a striped set of disks	Similar to RAID 1 and RAID 0 where you start with a striped set of disks and then mirror them.	Provides higher performance and fault tolerance than other RAID levels.	It requires more disks.

TIP To obtain more data on RAID, see https://datatracker.ietf.org/wg/raidmib/about/.

RAID 0 Striping

RAID 0 striping provides no fault tolerance or redundancy; the failure of one drive will cause the entire array to fail. As a result of having data striped across *x* number of disks, the failure will be a total loss of data. RAID 0 is normally used to increase performance and provide additional space (larger than a single drive) because RAID 0 appears as a single disk to the user. To implement RAID 0, you need to boot your computer to the BIOS and look for Disks and Settings for RAID (see Figure 13-1). There, you select the RAID level. In the example shown here, the ASUS motherboard has a built-in RAID controller like most modern motherboards do (see Figure 13-2). In this case, you select Create a RAID Volume, enter the name of the RAID set, select the RAID level, and select the disks to be part of the RAID set. Then you select Create Volume. Be warned that this will delete all data on all disks.

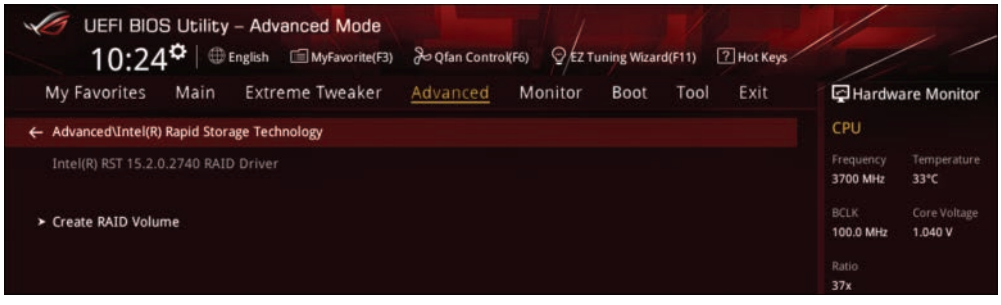


FIGURE 13-1 Bios RAID Controller

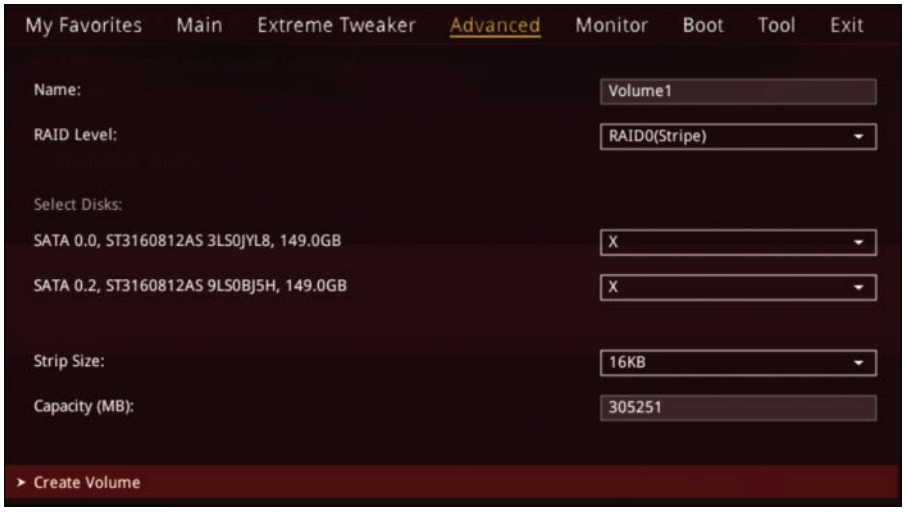


FIGURE 13-2 Creating a Volume

Your motherboard might be different. In fact, different ASUS models are slightly different in layout and configuration, but the steps are virtually the same.

RAID 1—Mirror

RAID 1—Mirror is exactly that—one or more disks with an additional disk to act as a mirror image for each disk. A classic RAID 1 mirrored disk pair contains two disks. The configuration offers no striping, no parity, and no spanning disks across multiple sets. Because all disks are mirrored, you need to match disk sizes to obtain the most benefit. The smallest disk in a mirror set is the largest the array can be. The performance of the RAID 1 array depends on the drive array, controller, slowest disk speed, and I/O load. Typically, you will see slightly slower speeds than a single drive. To configure RAID 1, you follow the steps described for Figure 13-1 and Figure 13-2, and select RAID 1 Mirror for the RAID level.

RAID 5—Striping with Parity

RAID 5—Striping with Parity consists of block-level striping with distributed parity among all drives. RAID 5 requires at least three disks and can suffer only a single drive failure before data becomes corrupted. The disk parity can be distributed in three different ways; only more expensive controllers allow you to select which method. The default sequence of data blocks is written left to right or right to left on the disk array of disks 0–X, also called Left Asynchronous RAID 5. Performance of RAID 5 is increased because all RAID members participate in the servicing of write requests. To configure and implement RAID 5, you follow the steps outlined for Figure 13-1 and Figure 13-2 and select RAID 5 Striping with Parity for the RAID level, of course taking into account your specific motherboard capabilities.

RAID 6—Striping with Double Parity

RAID 6—Striping with Double Parity consists of block-level striping with two parity blocks distributed across all member disks. Similar to RAID 5, there are different capabilities for how data is striped over the disk. None of them provide any greater performance over the other, and RAID 6 takes a performance penalty during read operations but is similar in the write speeds. You can suffer the loss of two disks using RAID 6 and still remain operational. To configure RAID 6, you follow the steps outlined for Figure 13-1 and Figure 13-2 and select RAID 6 for the RAID level. Also, note that higher-end server motherboards and controllers typically only have this RAID level capability.

RAID 10—Stripe and Mirror

RAID 10—Stripe and Mirror is the combination of creating a stripe and then mirroring it. RAID 10 is often called RAID 1 + 0 and array of mirrors, which may be two- or three-way mirrors. This configuration requires a minimum of four drives, and performance-wise it beats all RAID levels with better throughput and lower latency, except for RAID 0. To configure RAID 10, you need to consult your server manufacturer's guide on creating RAID 10 arrays.

Multipath

Key Topic

Storage ***multipath*** involves the establishment of multiple physical routes; for example, multipath I/O defines more than one physical path between the CPU in a computer and its mass-storage devices through the buses, controllers, and bridge devices connecting them. Storage area networks (SANs), covered later in this chapter, are primary examples of storage multipaths. When building a SAN, you create two paths from the server to the SAN switch, providing control from the SAN switch to the storage controller all the way to the disk.

Network Resilience

Network resilience involves more than just redundancy; it enables service operations in the face of outages, faults, and challenges. Threats and challenges for services can range from simple misconfiguration to large-scale natural disasters to targeted attacks. Network availability can be handled from multiple network cards, load balancers, and redundant network hardware. You can implement network resilience in multiple ways, as mentioned previously, such as having multiple network cards in each server, using load balancers to offload traffic to hosts that are less busy, and adding in layers of redundant networking components like switches and routers that are configured in high-availability configurations.

Key Topic

Load Balancers

There are several types of load balancers. Network load balancers automatically distribute incoming traffic across multiple network paths. This capability increases the availability of your network during component outages, misconfigurations, attacks, and hardware failures. Most network load balancers utilize routing protocols to help make path decisions during a network overload or network failure. A load balancer spreads out the network load to various switches, routers, and servers. Server load balancers distribute the traffic load based on specific selected algorithms to each of the servers in a group. Load balancers are typically placed in front of a group of servers to evenly distribute the load among them.

In network load balancing, when a router learns multiple routes to a specific network via routing protocols such as RIP, EIGRP, and OSPF, it installs the route with the lowest administrative distance in the routing table. If the router receives and installs multiple paths with the same administrative distance and cost to a destination, load balancing of network traffic will occur.

To mitigate risks associated with failures of load balancers themselves, you can deploy two servers in an active/active or active/passive configuration. In active/active, traffic is split between the two servers typically in a round-robin fashion. In active/passive configuration, all traffic is sent to the active server, and the passive server is automatically promoted to active status if the current active server fails or is taken down for maintenance.

Network Interface Card (NIC) Teaming

Network interface card (NIC) teaming allows you to group server network cards to increase throughput, improve performance, and provide path diversity and fault tolerance in the event of a single network adapter, cable, or path failure. NIC teaming is typically done via the system's motherboard/supervisor or via the cards themselves. Note that software NIC teaming is extremely inefficient and is generally not used any longer. NIC teaming allows a NIC to be grouped with multiple physical NICs, forming a logical network device known as a *bond*. This bond provides for fault tolerance and load balancing. To implement NIC teaming, typically the server motherboard with built-in NIC cards has a separate utility to configure the specific capability and mode.

Power Resilience

Power resilience is about ensuring a business, data center, or system has a reliable, regular supply of energy and contingency measures in place in the event of a power failure or fluctuations. Faults can include power surges, weather, natural disasters, accidents, and equipment failure.



Uninterruptible Power Supply (UPS)

An uninterruptible power supply or uninterruptible power source (UPS) is an electrical device that provides emergency power to a load when the input power source or main power fails. A UPS differs from an auxiliary or emergency power system or standby generator in that it will provide near-instantaneous protection from input power interruptions by supplying energy stored in batteries or supercapacitors. A UPS will allow you enough time to properly and safely power down equipment

until power is fully restored. A UPS has a limited run time due to the storage capability of batteries and is meant to be a temporary solution—in most cases, enough time to properly shut down resources. Some newer UPSs and servers are aware of each other, and when power fails, a signal is sent to the server asking it to properly shut down.

NOTE UPSs are used to protect electronic equipment and provide immediate emergency power in case of complete power failure.

Key Topic

Generators

A ***generator*** provides power to spaces and devices during complete power loss, blackouts, or in areas where standard electrical service isn't available. Generators range from 800 watts to over 500,000 watts, and there are different types of generators for every deployment requirement. Generators convert mechanical or chemical energy into electrical energy. They do this by capturing the power of motion and turning it into electrical energy by forcing electrons from the external source through an electrical circuit. The formula to figure out your required size generator is (Running wattage × 3) + Starting wattage = Total wattage needed. As mentioned previously, implementing a corporate generator strategy requires preplanning, sizing and capacity planning, and ensuring you build a sufficient platform to house the generators, fuel supply, and power transfer equipment.

Key Topic

Dual Supply

Dual supply is often confused with *dual power supplies*, where two power supplies are provided to a device, such as a server. For example, two power supplies with a single source provide power potential to the server, usually from a single "main circuit." If one power supply fails, the other power supply continues to provide power. Dual supply power is more complicated. Most facilities receive power from a single power company, single supply path, and single substation. For a substantial price, a company can request a design and buildout of a two-pathed independent power source solution. Dual supply power provides the ultimate in protection from power outages, and rarely fails. Figure 13-3 shows how dual supply power can provide resilience. Depending on the distance, one or two substations can provide power to the facility.

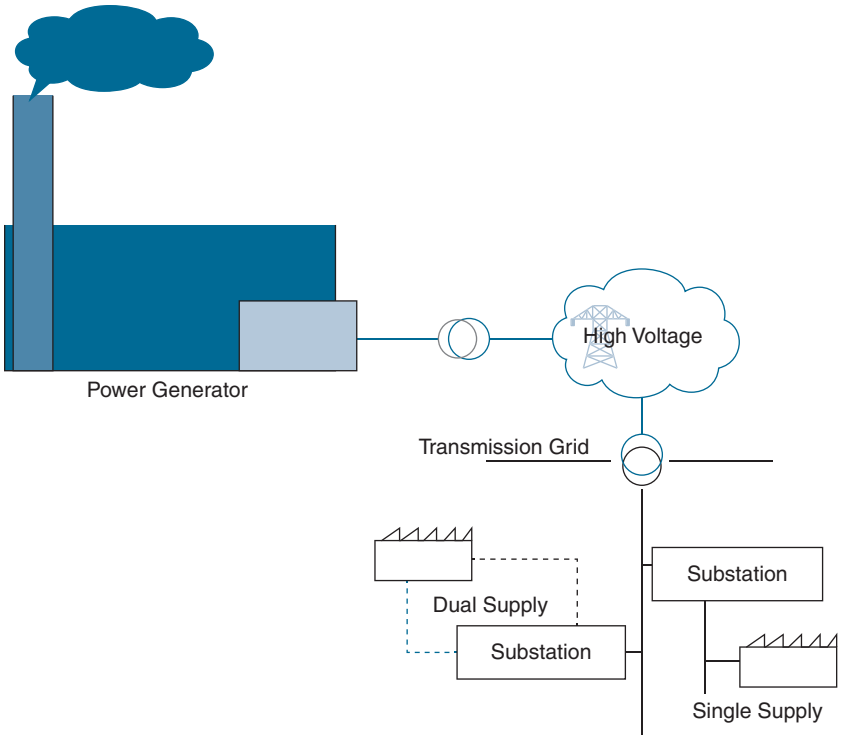


FIGURE 13-3 Resilience Through Dual Supply Power

Dual supply power mitigates the failure of a single substation or power feed that might be susceptible to a car hitting a light pole or a backhoe digging up a power cable.

Key Topic

Managed Power Distribution Units (PDUs)

A ***managed power distribution unit (PDU)*** is essentially a power strip with multiple outputs designed to distribute electric power to critical equipment. It is specifically designed to deliver power to racks of computers and networking equipment located within a data center. Certain PDUs have features that allow them to take power from two different power sources, known as A power and B power. Data centers face challenges in power consumption, protection, and management solutions, but PDUs provide monitoring capabilities to improve efficiency, uptime, and growth. There are two main categories of PDUs: basic PDUs and intelligent (networked) PDUs (or iPDUs). Intelligent PDUs normally have an intelligence module that allows them

to be remotely managed and monitored, such as the ability to view power metering information, power outlet on/off control, alarms, and sensors and to improve the power quality and provide load balancing.

Replication

When systems are not designed or configured with redundancy in mind, they create a single point of failure; nowadays, many components of the modern data center provide capabilities that allow for ***replication***, from storage area networks to virtual machines (VMs) that support those systems.

Key Topic

Storage Area Network

Storage area networks (SANs) are the most common storage networking architecture used by enterprises for business-critical applications that need to deliver high throughput and low latency. By storing data in centralized shared storage, SANs enable organizations to apply consistent methodologies and tools for security, data protection, and disaster recovery.

SANs are designed to remove single points of failure, making them highly available and resilient. A well-designed SAN can easily withstand multiple component or device failures. SAN-to-SAN replication (or synchronization) provides a fast no-impact high-performance backup of your all the data.

With replication, you could have a SAN in one data center; then you could send data to another SAN in a different data center location. You can back up very large volumes of data very quickly using SAN technology, and you can also transfer that data to a secondary location independently of your snapshot schedule. This approach is more efficient because traditional backup windows can take a very long time and impact the performance of your system. When everything is kept on the SAN, backups can be done very fast, and data can be copied in the background, so it's not impacting the performance of your systems. Figure 13-4 shows how a SAN accomplishes its high availability with dual paths.

In Figure 13-4, SAN-connected servers contain special fiber interface cards called host bus adapters (HBAs) that are configured as HBA1 and HBA2. The fiber is then connected to a SAN switch, which in turn connects to a SAN storage array. These two paths ensure diverse and high-availability paths to the server and the SAN storage.

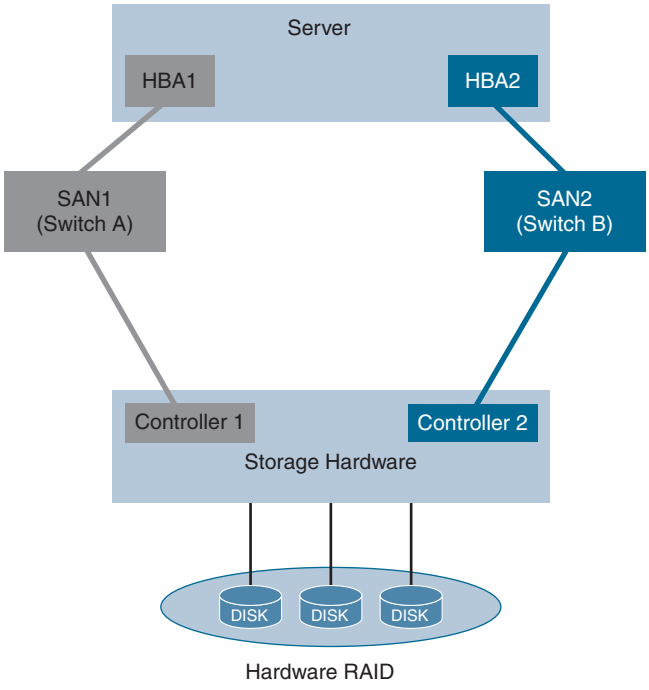


FIGURE 13-4 Storage Area Network

Key Topic

Virtual Machines

A ***virtual machine (VM)*** is a computer file, typically called an image, that behaves like an actual computer. It runs in a window, much like any other program, giving end users the same experience on a virtual machine as they would have on the host operating system itself. The virtual machine is sandboxed from the rest of the system, meaning that the software inside a virtual machine can't escape or tamper with the computer itself. Multiple virtual machines can run simultaneously on the same physical computer, typically a server. For servers, the multiple operating systems run side by side with a piece of software called a hypervisor to manage them.

Each virtual machine provides its own "abstracted" virtual hardware, including CPUs, memory, hard drives, network interfaces, and other devices. The virtual hardware is then mapped to the real hardware on the physical machine, which saves costs by reducing the need for physical hardware systems, along with the associated maintenance costs, and reduces power and cooling demand.

To ensure the availability of your virtual machines, you can enable replication. Several different types of VM replication give you different restore points, including a near real-time VM replication option that immediately copies the data as it is being

written to the VM, giving you a close-to-exact copy of the VM in near real-time. A more point-in-time replication option occurs on a scheduled basis. For example, if you replicate your VM once an hour, you can either have it automatically start when the protected VM fails or manually start when you want it to be turned on. Typically, you configure replication between data centers and the "vmotion," a virtual machine from one physical system to another system in the secondary data center.

TIP A hypervisor controls the virtual machines much like the supervisor of a switch: It controls distribution of the abstracted devices of a VM.



On-premises vs. Cloud

On-premises or cloud-based vulnerabilities will follow poorly configured and deployed assets from servers to networking equipment. If shortcuts are taken and best practices are not followed, your company may experience a cyber attack that will affect your equipment no matter where it lives. *On-premises* and *cloud-based* are simply terms that describe where systems store data. Many of the same vulnerabilities that affect on-premises systems also affect cloud-based systems.

On-premises redundancy provides rapid recovery from small and localized events, such as a server failure, rack failure, or single-source power failure. However, when more catastrophic events occur, on-premises recovery can take an extremely long time. Locating and replacing damaged equipment and locating operational space to stage a recovery add to the delayed recovery of the company's business operations.

Cloud-based redundancy allows the enterprise to recover in a moderate to rapid manner, depending on the deployment of resources. For example, if the company's users are all stored or managed in a cloud-based directory service like Azure Active Directory (AD), recovery of users and rights to local assets can take place more quickly. If the company's primary assets are mostly cloud based, such as email, databases, and enterprise resource planning (ERP), then the outage is rather confined. In this case, the company need only move to a suitable facility and connect its users to the cloud or have employees work from home.

Cloud-based failures could extend the company's recovery times by days, if not weeks; therefore, it is prudent for the company to regularly back up its cloud-based assets. But there is never any guarantee that a specific cloud provider will be around tomorrow.

Backup Types

There is nothing more important to an organization than its data. Without it, most companies would cease to exist. If a company does not take proper precautions to back up its data and that data is damaged or destroyed, the company may no longer be able to operate. Backing up critical data is one of the highest priorities a company can undertake. Table 13-3 displays the backup types and how they provide protection. The most common types are full backup, differential, and incremental. A company only having backups in the cloud puts those backups at risk if the provider has geographical outages or goes out of business; or if the provider gets compromised, the backups could be transferred to the attacker and ransomed or deleted/destroyed.

Key Topic

Table 13-3 Backup Type Benefits and Drawbacks

Type	Benefits	Drawbacks
Full	Provides full copy of data set. Provides the best protection.	Time-consuming to restore. Requires lots of storage.
Incremental	Takes less time and storage space.	Time-consuming to restore. Requires all backups in the backup chain to restore.
Snapshot	Creates fast backups of specific times.	May take more space to restore each snapshot in order.
Differential	Takes a shorter restore time than incremental.	Can grow to be much bigger in size than incremental.
Tape	Offers the best protection for catastrophic disasters if proper rotation and offsite storage are employed.	Is costly, very slow to recover and restore; requires specific order.
Disk	Provides one of the fastest and most complete backup and restore capabilities.	Is costly.
Copy	Provides quick backup.	Doesn't copy in-use files; not complete.
NAS (network-attached storage)	Is PC connected or network connected and shared among many; provides readily available storage.	Is very slow, reliant on the connected PC and/or the network speed and availability.
Cloud	Provides connected backup storage anywhere in the world with private/public cloud storage.	Is slow to back up, slow to restore, and dependent on network availability.
Image	Provides the most complete type of backup, including the entire operating system and all other files. Can be directly connected, NAS, SAN, or cloud based.	Requires a large amount of storage.
Online vs. offline	Online storage is readily available, whereas offline storage provides safety from potential hackers and viruses.	Online backups and storage are costly. Offline storage requires physically connecting or turning on specific devices.

Type	Benefits	Drawbacks
Offsite storage	Is typically rotated because tapes are stored in an offsite facility.	Can delay the recovery process.
Distance considerations	Provides geographical diversity if backed up over long distances.	Can take a long time to recover backups that are located geographically far away, depending on bandwidth and facility availability.

When you're backing up data to the cloud, a large concern is who owns the data, what privacy laws and regulations apply, and whether data is stored across borders.

Figure 13-5 illustrates the most common backup types, which are further described in the sections that follow.

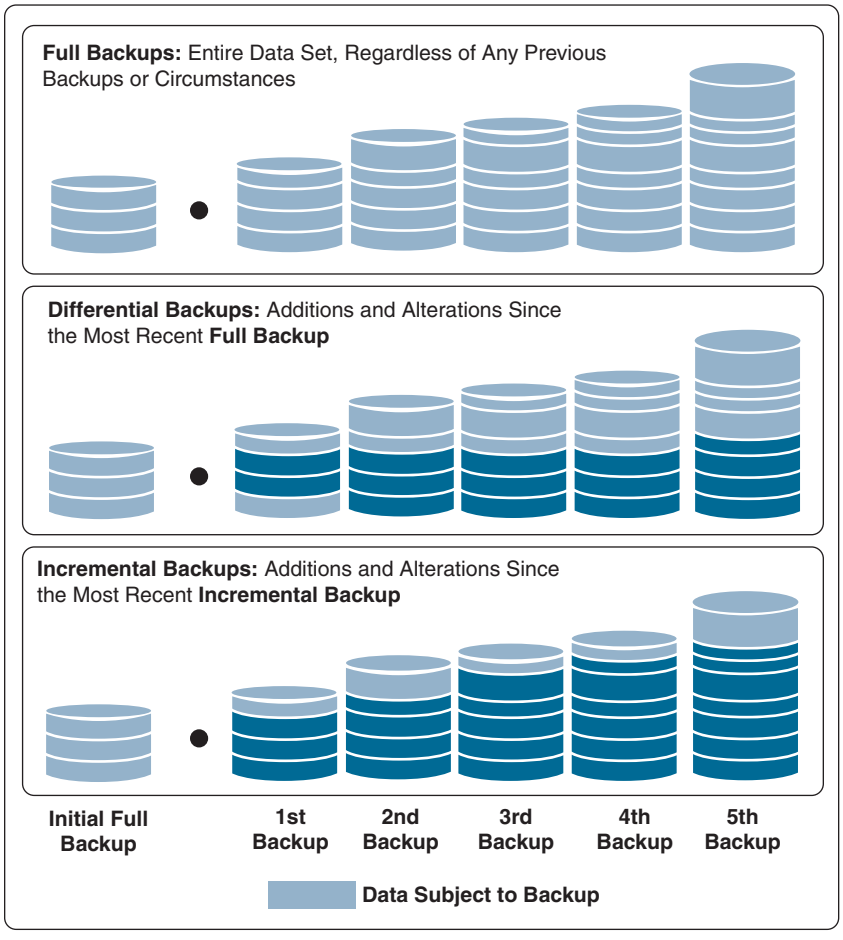


FIGURE 13-5 Backup Types

Full Backup

A full backup is the most complete type of backup where you make a complete copy of all the selected data. This includes files, folders, applications, hard drives, and more. The highlight of a full backup is the minimal time required to restore data; however, because everything is backed up at one time, it takes longer to back up compared to other types of backups. Most businesses tend to run a full backup and occasionally follow it up with a differential or incremental backup. This approach reduces the burden on the storage space, increasing backup speed.

Differential Backup

A differential backup is in between a full and an incremental backup. This type of backup involves backing up data that was created or changed since the last full backup. A full backup is done initially, and then subsequent backups are run to include all the changes made to the files and folders. It lets you restore data faster than a full backup because it requires only two backup components: the initial full backup and the latest differential backup.

A differential backup works like this:

Day 1: Schedule a full backup.

Day 2: Schedule a differential backup. It covers all the changes that took place between Day 1 and Day 2.

Day 3: Schedule a differential backup. It makes a copy of all the data that has changed from Day 2 (this includes the full backup on Day 1 + differential backup) and Day 3.

Incremental Backup

The full backup is first; then succeeding backups store only changes that were made to the previous backup. An incremental backup requires space to store only the changes (increments), which enables fast backups.

Non-persistence

Non-persistent data is that which is not available after fully closing the application or turning off the system. Non-persistent desktops have a short existence under natural conditions. They are known as a many-to-one ratio, meaning that they are shared among end users, but no changes are saved. Non-persistent storage is typically short-lived storage used for logs and diagnostics. Non-persistent capabilities, such as those in bootable distributions, allow you to boot up a full operating system

and perform tasks, browse, and operate, and then when the machine is powered down, everything changed or modified is gone.

Key Topic

- In ***revert to known state*** the desired outcome is to take a system back to a prior moment in time or state of existence. In the process, typically changes implemented since that known "good" point in time are removed, usually to start at a point and time before a problem or issue arose; the purpose is to reset a device to a stable and secure setting. Revert to known state examples include Cisco's Configuration Rollback, VMware snapshot, Windows Steady State, Linux Systemback, and Mac's Time Machine.
- ***Last known good configuration***, or LKGC for short, returns the configuration of a system back to just before a problem started. You can undo recent changes that caused a problem or error or that weakened security. In Windows 7 and 8 computers, you can fix driver and configuration issues by selecting F8. In Windows 10, Microsoft disabled this capability; however, you can still enter Safe Mode (a last good configuration alternative) in the Windows Recovery Advanced Options.

> **NOTE** For more information, see https://softwarekeep.com/help-center/how-to-boot-windows-10-into-the-last-known-good-configuration.

- ***Live boot media*** is a lightweight bootable image on a USB flash drive or external hard disk drive containing a full operating system that can be booted to. The added benefit of writable storage is that it allows customizations to the booted operating system. Live USBs can be used in embedded systems for system administration, data recovery, or test driving an operating system, and can persistently save settings and install software packages on the USB device.

Key Topic

High Availability

High availability (HA) is a characteristic of a system that aims to ensure an agreed-upon level of operational performance, usually uptime, for a higher-than-normal period. Modernization has resulted in an increased reliance on these systems, hence the need to make them available at all times. Availability refers to the ability of users to obtain a service or good or to access the system—whether to submit new work, update or alter existing work, or collect the results of previous work. Three principles of system design in reliability engineering can help achieve high availability:

- **Elimination of single points of failure:** This means adding or building redundancy into the system so that failure of a component does not mean failure of the entire system.

- **Reliable crossover:** In redundant systems, the crossover point itself tends to become a single point of failure. Reliable systems must provide for reliable crossover.
- **Detection of failures as they occur:** If the preceding two principles are observed, a user may never see a failure, but the maintenance activity is a must.

High availability is measured as the available time of a system being reachable. Assuming 100 percent is never failing, HA is the widely held but difficult-to-achieve standard of availability of a system or service known as "five nines," or 99.999 percent available. Availability percentage is measured as follows:

Availability = (minutes in a month – minutes of downtime) * 100/minutes in a month

For example, if there is a service-level agreement for five nines, the end user can expect that the service will be unavailable *for 26.3 seconds per month and 5 minutes and 15 seconds per year*.

Scalability is the property of a system to handle a growing amount of work by adding resources to the system. In terms of systems, it must be able to scale to number of users and workloads.

Restoration Order

Prior to a disaster, a company should run through an exercise called disaster recovery planning (DRP). From here, you can build a plan that will include all systems you plan to restore during the recovery process and in which order. Certain systems are required to be online prior to anything else; typically, they are networking and authentication services, like the physical network connectivity, and Active Directory. You can build a table similar to the sample shown in Table 13-4 to ensure everyone on the team is on the same page. That way, any replacements or new resources who were not involved in DRP planning or testing will be able to understand how things come together.

Key Topic

Table 13-4 Disaster Recovery Order

Process	Priority	System
Connectivity	Restoring specific networks that include administrative networks.	Routers, switches, and networking infrastructure; Internet; telco services.
Authentication	Performing logical layout of systems and interdependencies.	Active Directory and authentication services.
Phone systems	Restoring phones (a very high priority).	Requirements for emergency response, health and safety, E911 (Enhanced 911), and support.

Process	Priority	System
CRM/SAP/HRM	Restoring systems that enable the corporation to perform operations and provide continuity.	Ordering systems, online sales and marketing systems, human resources.
Other resources	Continue to build this out.	List all your important systems and put them in order.

Key Topic

Diversity

One of the best ways to ensure availability is through ***diversity*** and redundancy. The more diverse the redundancy, the more fault tolerant the systems and network. You can minimize the business impact of a cyber attack on your organization by providing diversity in the manner you build and deploy your network and systems.

Technologies

Diversity with technology starts with utilizing different methods to accomplish similar tasks. For example, you can provide network resilience by utilizing a routing protocol and a load balancer. Although these are two completely different methods to provide resilience, they provide diversity that helps a company recover from a failure. By adding two network cards to a server, you can utilize diverse paths.

Vendors

Diversity in vendors provides resilience in your supply chain, and having relationships with multiple vendors ensures that your procurement process won't break down if one of your vendors experiences an attack or disaster, or it goes out of business. Employing multiple vendors allows you to obtain the best price when sourcing but also ensures access to products when they may be in short supply.

Crypto

Cryptography (encryption) should be utilized as often as possible. For example, your backups should be encrypted with a strong encryption method. This way, if they are stolen, the attacker will not have access to the data. Windows servers and workstations can encrypt hard disks (storage). Any link that leaves your organization should be encrypted to ensure your communications are private. Encryption should be utilized and implemented with every application and system your organization uses.

Controls

The control systems in your building include the HVAC system, elevators, lighting, and other building controls you take for granted while at work. Control systems are purposely designed with diversity and high availability in mind so that you aren't stuck on an elevator, do not lose lighting, or are unable to exit a building during a power failure. These feats are accomplished by running dual-diverse path cables to each component, having redundant controllers, and using automation so that if power does fail, the elevator is automatically taken to the ground floor with the doors open and locked.

TIP Control systems are unique in that they manage and monitor mission-critical systems such as lights, power, generators, and data centers.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 13-5 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 13-5 Key Topics for Chapter 13

Key Topic Element	Description	Page Number
Paragraph	Redundancy through geographic dispersal	315
Section	Redundant Array of Inexpensive Disks	316
Paragraph	Redundancy through storage multipath	319
Section	Load Balancers	319
Section	Uninterruptible Power Supply (UPS)	320
Section	Generators	321
Section	Dual Supply	321
Section	Managed Power Distribution Units (PDUs)	322
Section	Storage Area Network	323
Section	Virtual Machines	324

Key Topic Element	Description	Page Number
Section	On-premises vs. Cloud	325
Table 13-3	Backup Type Benefits and Drawbacks	326
List	Non-persistent capabilities	329
Section	High Availability	329
Table 13-4	Disaster Recovery Order	330
Section	Diversity	331

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

geographical dispersal, disk redundancy, network redundancy, Redundant Array of Inexpensive Disks, multipath, network interface card (NIC) teaming, generator, dual supply, managed power distribution unit (PDU), replication, storage area networks (SANs), virtual machine (VM), full backups, incremental backups, snapshot backups, differential backups, tape backups, disk backups, copy backups, NAS (network-attached storage) backups, cloud backups, image backups, online vs. offline backups, offsite storage backups, revert to known state, last known-good configuration, live boot media, high availability, scalability, diversity

Review Questions

1. What is geographical dispersal?
2. What is disk redundancy?
3. What is a UPS?
4. What is replication?
5. What is reverting to known state?



This chapter covers the following topics related to Objective 2.6 (Explain the security implications of embedded and specialized systems) of the CompTIA Security+ SY0-601 certification exam:

- Embedded systems
  - Raspberry Pi
  - Field-programmable gate array (FPGA)
  - Arduino
- Supervisor control and data acquisition (SCADA)/Industrial control system (ICS)
  - Facilities
  - Industrial
  - Manufacturing
  - Energy
  - Logistics
- Internet of Things (IoT)
  - Sensors
  - Smart devices
  - Wearables
  - Facility automation
  - Weak defaults
- Specialized
  - Medical systems
  - Vehicles
  - Aircraft
  - Smart meters
- Voice over IP (VoIP)
- Heating, ventilation, air conditioning (HVAC)
- Drones
- Multifunction printer (MFP)
- Real-time operating system (RTOS)
- Surveillance systems
- System on chip (SoC)
- Communication Considerations
  - 5G
  - Narrow-band
  - Baseband radio
  - Subscriber identity module (SIM) cards
  - Zigbee
- Constraints
  - Power
  - Compute
  - Network
  - Crypto
  - Inability to patch
  - Authentication
  - Range
  - Cost
  - Implied Trust

Understanding the Security Implications of Embedded and Specialized Systems

As we are striving to fit more functionalities into smaller embedded systems, their security is often neglected. Internet of Things (IoT) mobile devices such as smartphones, tablets, and wearables are already ubiquitous, but as the IoT market is expanding, nonmobile IoT systems are expected to outnumber the current IoT-enabled mobile devices in a few years. Embedded systems are getting smaller and smarter, enabling us to get more things done than ever before. As more functionalities are embedded in smaller device footprints, there also is an upsurge in security concerns. Device vendors prefer to add new features that often crowd out the basic security systems, thus launching devices with low security capabilities.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 14-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 14-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Embedded Systems	1–2
SCADA/ICS	3
Internet of Things (IoT)	4–5
Specialized Systems	6
Surveillance Systems	7
Communication Considerations	8

Foundation Topics Section	Questions
Embedded System Constraints	9
System on a Chip (SoC)	10
Real-Time Operating Systems (RTOS)	11
Drones	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following embedded devices are field programmable and are a reconfigurable integrated circuit that can perform logic gates such as AND gate, OR gate, and NOT gate?
   a. FPGA
   b. PStat
   c. BIOS
   d. None of these answers are correct.

2. Which of the following Supervisory Control and Data Acquisition systems are part of the logistics ecosystem?
   a. Just-in-time manufacturing
   b. Fish tank controls
   c. Industrial plant mixing platforms
   d. None of these answers are correct.

3. Which of the following embedded devices can read plant field device output and input and can control or regulate devices like motors and generators?
   a. PLC
   b. PStat
   c. Bios
   d. None of these answers are correct.

4. Which IoT device is considered a wearable device and is helpful for your well-being?
   a. Smart tennis shoes
   b. Smart watch
   c. Smart jackets
   d. Smart headband

5. Which of the following IoT devices are part of facility automation?
   a. Lighting and air conditioning
   b. Raspberry Pi
   c. Garage door sensors
   d. Computer Energy Star

6. Security associated with specialized systems generally requires the user to set up logging and auditing of passwords. What else should the user do to ensure the system is secure?
   a. Lock passwords in a password vault that only staff has access to.
   b. Ensure the network operations center is aware of hard drive failures in critical systems.
   c. Always log passwords in the security operations center log book.
   d. Always create and rotate very complex passwords.

7. Today's surveillance systems are capable of using technology to perform all of the tasks listed here except for which of the following?
   a. Monitoring the data center access
   b. Pivoting from one system to another
   c. Performing facial recognition
   d. Authenticating application usage

8. Which communication method is part of the IEEE's 802.15.4 standard and has a transmission rate of 250 kbps?
   a. 5G
   b. Narrow band
   c. Baseband radio
   d. Zigbee

9. A major constraint on embedded systems deals with the size of the unit and capability. What is one of the concerns related to security?
    a. The amount of energy required to load the boot loader at startup
    b. The ability to reboot the system remotely with a password
    c. The amount of code required to implement a secure algorithm
    d. The amount of power and amperage available to power on the unit at startup

10. Which of the following is considered a system on a chip?
    a. The desktop computer used to perform number calculations that can be used to create documents
    b. The ability to reboot the system remotely with a password
    c. A cell phone and cell phone–type device with personal data access
    d. A tablet with WiFi and/or cell LTE connectivity capabilities

11. Which is one application of a real-time operating system being used in the commercial space that helps save lives?
    a. Nucleus RTOS from Mentor, used in medical devices
    b. Raspberry PI using the optional POE module for power
    c. A SOC that can be used to provide fusion
    d. The Arduino med-xcapp heart register system connected to a WiFi network

12. Which low-power device has been used in commercial applications to reduce human interaction capabilities and keep people safe?
    a. Catapults with wireless controllers that can move objects across enemy lines and determine troop sizes
    b. Drones used in harsh environments such as refineries and terminals where flammable gases can be dangerous to humans
    c. Robotic window washers, which ensure that windows are washed without the need for a human to be on a scaffolding
    d. Microdots that enable spies to see inside facilities to ensure it's safe for them to carry out operations

Foundation Topics

Embedded Systems

Embedded systems are microprocessor-based computer hardware systems with software that is designed to perform a dedicated function, either as an independent system or as a part of a larger system. At the core is an integrated circuit designed to carry out computation for real-time operations. The following embedded systems are used in everyday projects, systems, and commercial products:

Key Topic

- ***Raspberry Pi*:** A low-cost, credit-card-sized computer (see Figure 14-1) that plugs into a computer monitor, keyboard, and mouse. It's capable of doing everything a desktop computer can do—from browsing the Internet and playing high-definition video to creating spreadsheets, doing word processing, and playing games. If you watched *Mr. Robot*, Season 1, Episode 4, you saw Elliott use a Raspberry Pi to hack the HVAC system to raise the temperature to destroy backup tapes.

> **NOTE** The Raspberry Pi is a capable computer that can interact with the outside world and can be used in a wide array of digital maker projects, from hacking projects, music machines, and weather stations to tweeting birdhouses with infrared cameras.

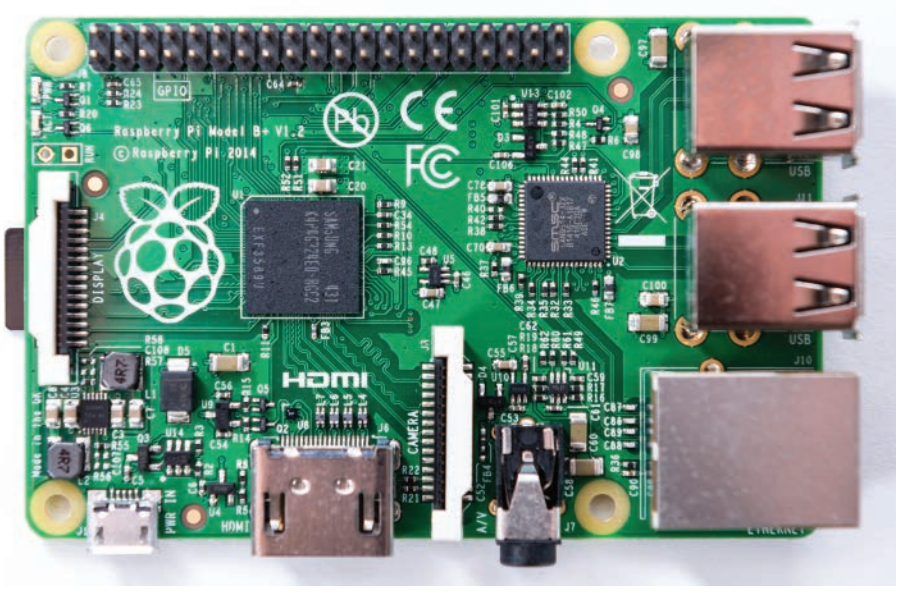


FIGURE 14-1 Raspberry Pi

NOTE For more information on Raspberry Pi, visit www.raspberrypi.org/.

- ***Arduino*:** An open-source electronics platform based on easy-to-use hardware and software. Arduino boards are extremely flexible; they are able to read inputs such as a light on a sensor, button presses, and even a Twitter message (and turn it into an output), activating a motor, turning on an LED, and publishing something online. You can tell your board what to do by sending a set of instructions. The Arduino software is easy for beginners to use yet flexible enough for advanced users. It runs on Mac, Windows, and Linux. Teachers and students use it to build low-cost scientific instruments, to prove chemistry and physics principles, and to get started with programming and robotics. Designers and architects build interactive prototypes, and musicians and artists use it for installations and to experiment with new musical instruments. Arduino is like any compute platform, so it is susceptible to a series of attacks including buffer overflows, stack smashing (data leakage), and code execution. These devices also run Linux operating systems, which also can be exploited. Care should be taken where these devices run, protecting them from exposure to unprotected networks and systems.

TIP You can visit the Arduino open-source community at https://arduino.cc for all kinds of projects and Arduino device add-ons.

- ***Field-Programmable Gate Array (FPGA)*:** An integrated circuit designed to be configured by a customer or designer after manufacturing—hence the term *field-programmable*. Modern FPGAs consist of mixes of configurable static random-access memory (SRAM), high-speed input/output (I/O) pins, logic blocks, and routing. More specifically, an FPGA contains programmable logic elements called logic elements (LEs), as well as a hierarchy of reconfigurable interconnects that allow the LEs to be physically connected to one another. You can configure LEs to do complex functions or simply perform basic logic gates, such as AND gate, OR gate, and NOT gate. FPGAs are vulnerable to supply chain attacks because chips are manufactured in a number of countries with varying degrees of oversight and attention to processes. FPGAs are particularly vulnerable to readback attacks directly obtained when the device is in the Field mode. There are killswitch and backdoor attacks in which the device itself can be programmed with specific code, waiting for an attacker to initiate it remotely.

Key Topic Supervisory Control and Data Acquisition (SCADA)/Industrial Control Systems (ICS)

Supervisory control and data acquisition (SCADA) systems and ***industrial control systems (ICSs)*** are made up of a combination of specialized software and hardware elements that allow organizations to

- Monitor equipment processes and performance
- Perform data monitoring and collection
- Facilitate data analysis and storage of analyzed data
- Perform remote control of equipment, devices, components, and processes
- Trigger alerts and alarms when performance or deviation aberrations are detected
- Report generators for scheduled and automated variance reporting and comparison, as well as production efficiency and status code statistics

SCADA systems are crucial for industrial organizations because they help to maintain efficiency, process data for smarter decisions, and communicate system issues to help mitigate downtime.

The basic SCADA architecture starts with a series of configurable devices called programmable logic controllers (PLCs) or remote terminal units (RTUs). PLCs and RTUs are microcomputers that communicate with an array of objects such as turbines, machines, gauges, sensors, and end devices, and then route the information from those objects to computers with SCADA software. The SCADA software processes, distributes, and displays the data, helping operators and other employees analyze the data and make important decisions.

For example, the SCADA system quickly notifies an operator that a batch of products shows a high incidence of errors. The operator pauses the operation and views the SCADA system data via a human-machine interface (HMI) to determine the cause of the issue. The operator reviews the data and discovers that a specific machine in the process was malfunctioning. The SCADA system's capability to notify the operator of an issue helps him or her to resolve it and prevent further loss.

To defend against cyber attacks, a firewall can act as a security buffer between the SCADA network and the corporate network. SCADA systems should never be connected to the Internet, even with firewalls in place. You should configure internal firewalls that include virtual private networks (VPNs) for remote SCADA systems internal administration as an added extra layer of protection. To address

authentication and authorization security concerns, you need to configure the authentication process in the software to identify authorized users for access to the SCADA system, enable logging to track users' activity and alerting to warn of deviations to critical controls, and ensure least privilege access. Depending on the perceived risk of cyber attack on the SCADA system, you should deploy intrusion detection and offsite redundancy, as part of a defense-in-depth strategy:

- **Facilities:** From facility controls like HVAC, elevators, lighting, room thermal and temperature, occupancy, parking sensors, device usage sensors and much more, facility managers are able to monitor and control whether a building is operating optimally. SCADA systems can help organizations reduce their carbon footprint through energy management, by automatically turning off lights and air conditioning to sections of the building that are unoccupied.
- **Industrial:** ICS and SCADA systems monitor, control, sense, and warn engineers of all aspects of the processes in industrial plants. From refineries, chemical plants, and more, SCADA systems can control the flow of crude oil being turned into gasoline, asphalt, and chemicals like cleaning supplies. This complex process has many controls that could be dangerous if inappropriately configured.
- **Manufacturing:** Control systems in the manufacturing process depend heavily on data, performance monitoring, and capabilities to control remote equipment, which are among the most-discussed concerns in manufacturing. These components are frequently mentioned as the backbones of "Industry 4.0." Industry 4.0 is the current trend of automation and data exchange in manufacturing technologies. It includes cyber-physical systems, the Internet of Things, and cloud computing. Industry 4.0 creates what has been called the "smart factory." Monitoring, data capture and analysis, and equipment control from anywhere are recognized as highly effective ways to improve facility and equipment productivity, performance, and efficiency. The benefits of using SCADA control systems in the manufacturing process include reducing product errors and discards (thanks to earlier problem detection and remedies), improving productivity, and maximizing the effectiveness of machine uptime. Maintenance benefits through the use of data and sensors include more accurate diagnostics, more efficient maintenance processes, and the enablement of predictive maintenance. Data-driven insights into equipment performance can facilitate more accurate calibration, operation, and decision-making processes.
- **Energy:** SCADA systems aid in energy management because the global need for cost-effective electrical power is on the rise. With an increase in factory

automation productivity and accuracy, with ever-expanding data centers, there is a greater demand on electrical power utilities. An organization can integrate alternative energy sources, such as solar and wind, in their overall facility power source planning to obtain efficiency, stability, and safety. For power distribution system generators, operators and manufacturers alike, unplanned downtime is extremely costly, risky, and inconvenient. For operators, a stable, sustainable load flow across all components of a system is key for uninterrupted, safe operation. For manufacturers, any unplanned downtime can mean delays in production and missed deadlines. In both cases, predictive and proactive maintenance strategies can maximize uptime by alerting operators and manufacturers to potential issues early. Transit authorities use SCADA technology to regulate electricity to subways, trams, and trolley buses.

- ***Logistics*:** SCADA systems are capable of managing parts inventories for just-in-time manufacturing, regulating industrial automation and robots, and monitoring processes and quality control. Transit authorities use SCADA technology to automate traffic signals for rail systems; to track and locate trains, trucks, and buses; and to control railroad crossing gates. Continuous communication for trucks, locomotives, and transport vehicles provides valuable information, such as location and status, and is especially important for unmanned vehicles. SCADA systems can regulate traffic lights, control traffic flow, and detect out-of-order signals. Other specialized systems extend GPS and cargo tracking systems because knowing where ships, trains, trucks, and inventory are located anywhere in the world is critical to meet today's demand for products. These systems require special security considerations; regulators have long insisted that these systems be air-gapped to reduce threats. An air-gapped system is standalone in terms of a system on its own or a network that is isolated from the outside world. A key mitigation or control against attacks on SCADA systems and ICS is to implement physical segregation of internal and external networks, to reduce the attack surface by segregating the SCADA network from the corporate LAN with a firewall. An ICS is managed via a SCADA system. SCADA systems drive HMI for operators to monitor the status of a system while in the field. Other ICS components include industrial automation and control systems (IACS), distributed control systems (DCS), programmable logic controllers (PLCs), and remote terminal units (RTUs).

Figure 14-2 shows a basic SCADA diagram.

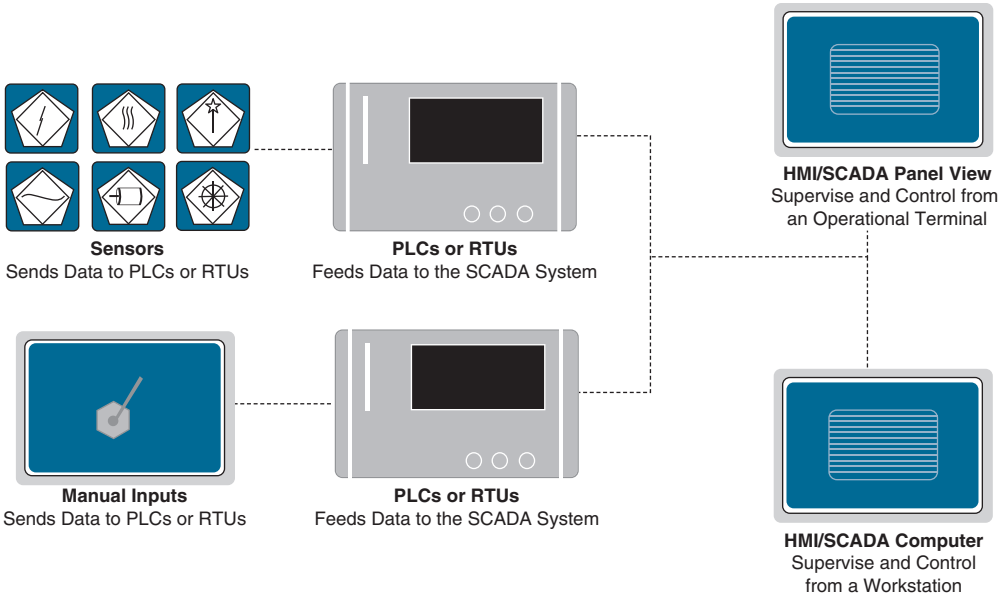


FIGURE 14-2 SCADA Diagram

Internet of Things (IoT)

The benefits of the ***Internet of Things (IoT)*** are undeniable; and yet high-profile attacks, combined with uncertainty about security best practices and their associated costs, are keeping many businesses from adopting the technology. Besides, end users are wary of the consequences of IoT security breaches. Recent research indicates that 90 percent of consumers lack confidence in IoT device security.

NOTE IoT manufacturers should do a thorough security risk assessment that examines vulnerabilities in devices and network systems as well as in user and customer back-end systems. Risk must be mitigated for the entire IoT lifecycle of the deployment, especially as it scales.

TIP You should change default manufacturer usernames and passwords.

Key Topic

Cybersecurity must be designed into IoT devices from the ground up and at all points in the ecosystem to prevent vulnerabilities in one part from jeopardizing the security of the entire system. Today's IoT ecosystems are complex; devices in virtually every industry, which includes cellular networks to cloud applications and

back ends, can be connected and configured to send data. Hackers are consistently looking for ways to exploit IoT system vulnerabilities. Because IoT devices are so diverse, there is no "one size fits all" cybersecurity solution that can protect any IoT deployment.

The following components are associated with the Internet of Things (IoT).

- ***Sensors*:** Sensor cybersecurity implications go well beyond the voice-controlled devices you talk to at home, such as your lights and TV. Sensors are part of the critical infrastructure that monitors and manages power plants, refineries, maritime ships, and air traffic controllers. Because sensors are small and versatile, functionality is crucial, but security is usually secondary and largely dependent on outside security systems. To protect sensors from attacks, you should use a layered and air-gapped approach when designing and deploying supporting networks.
- ***Smart devices*:** Smart devices are everywhere, from wired and wireless networks that control lighting to A/C, garage doors, security systems, and curtains; they are embedded in TVs, toasters, and even your barbecue grill. To protect smart devices, you should disable features that you are not using and enable them only when required, use strong encryption methods, and employ complex passwords. You should turn privacy settings to their highest levels and select to not share data, or not turn them on at all, and where possible, change the default manufacturer usernames and passwords.
- ***Wearables*:** Wearables include smart watches, pedometers, glasses, 3D headsets, smart clothing, health and fitness devices, heart and chemical and gas sensors, and much more. Security for these devices has been improving with the advent of high-performing processors but is still lacking in many aspects. We have seen the breach of smart watches that disclose running patterns. To protect wearables, much like most other devices, you should change all default settings, disable features and functions you're not using, set privacy settings to their highest level, and turn off location tracking where possible.
- ***Facility automation*:** An unprotected building automation system can quickly become a high risk with potential life-threatening damage to its occupants and the building itself; therefore, securing these systems is a necessity. Even the most severe data breach in a traditional organization is unlikely to have a direct physical impact on the life of occupants of the building. However, a malicious attack on a building automation system can very easily pose a significant threat to the health and safety of the occupants. Protecting these systems typically starts with a separate and isolated network from the corporate network, separate firewalls and monitoring, and abiding by a strict zero-trust approach to any connections made to this network.

- **Weak defaults:** From the implementation of weak protocols to default usernames and passwords, many of today's IoT devices are shipped with weak defaults. There are countless IoT devices and applications in your home and business; using weak or default passwords places devices at greater risk of exposing personal information or being used as pivot devices to gain a strong foothold inside. Pending legislation would require IoT device manufacturers to use strong encryption methods and require complex passwords during setup. However, until that happens, you must change default passwords and utilize the strongest encryption method the device supports. If you are unable to get past this, consider changing the device to something that has stronger capabilities. Don't use *Admin* as a username because it's the most common username used in attacks. Also, be sure to change default passwords to complex nondictionary types. In addition, you should run regular audits and scans on all IoT devices (several open-source IOT scanners are available), check for new and/or unknown devices on your network, and check the logs of those devices for password lockouts and strange activity. Another option is to check Shodan (a repository of Internet-accessible devices); this is where attackers go to find insecure devices to attack. Make sure your devices are not listed, but if they are, plug the device in behind a firewall, turn off unused features, and use complex secure passwords.

Specialized Systems

Embedded specialized systems like those used in hospitals, vehicles, aircraft, and smart meters have a long history of placing functionality and convenience before security. There are some simple and not so simple steps an organization can take to reduce the exposure that these devices present. First, you should always utilize the strongest encryption method available on the device. Second, always create and rotate very complex passwords. Third, set up a continuous method to log and audit access to the devices. Fourth, ensure you use the latest manufacturer updates to keep your devices up to date, and do not use beta, test, or bleeding edge updates. If all else fails, disable communication protocols such as WiFi, Bluetooth, Zigbee, and any others.

Specialized systems can include those systems that monitor other systems. For example, a supply chain attack was recently perpetrated by a specific nation outside the U.S.; we call these attacks Nation State attacks, and this specific attack perpetrator was nicknamed APT29 or Cozy Bear. The attack was against a rather well-known network and system monitoring product. Thousands of customers from the U.S. government to cybersecurity companies were affected. In the end, root cause analysis determined an intern used a weak password to protect a rather high-profile system used for pushing updates of its products to customers. (Ref: CVE-2020-166).

Another example of a high-profile attack comes from CVE-2020-14750—a remotely exploitable without authentication attack on a webserver technology, where an attacker can take complete control of the system. This specific web server was prevalent and deployed and used widely. These types of vulnerabilities are reported infrequently; however, when they hit, they affect large populations of systems.

As a security professional, you should always be aware and on alert for notifications and warnings of vulnerabilities to your environment and sign up for alerts from cve.mitre.org, us-cert.cisa.gov, and Cisco Talos to help stay ahead of attackers.



Medical Systems

The Linux operating system has become one of the main operating systems for embedded medical devices. Certain real-time operating systems (RTOSs) such as the Nucleus RTOS from Mentor can be acquired pre-certified, as can other embedded software components from a number of vendors for medical devices. Linux is generally handled using a concept from IEC 62304 called Software of Unknown Provenance (SOUP) for today's medical devices. Under SOUP guidelines, Linux is considered part of the risk assessment of the overall device, and potential failures of Linux as used in the device must be considered and mitigated if they might cause harm to a patient. This risk assessment must meet the requirements of the FDA's pre- and post-submission guidance. When it comes to safety, things that could go wrong will impact a patient's health. For example, you have likely seen movies in which someone becomes able to control a heart device (such as a pacemaker) and cause the death of an enemy. Consider a virus attacking the network controlling the updates to those devices. A manufacturer of insulin pumps did a complete recall due to a MITM replay attack that allowed hackers to replay traffic, essentially doubling someone's insulin. Medical device users should disable WiFi, Bluetooth, and other outside communications unless absolutely necessary and protect the connection with a strong wireless password and access control rules to limit communications to only vendor-specific IP address/domains.



Vehicles

Most new vehicles are equipped with some sort of connected infotainment system, such as Apple CarPlay, Android Auto, or a similar mobile device–connected standard, allowing for seamless integration with the operator's smartphone or tablet. In some cases, malware stored on a mobile device can be transferred to the automotive computer when connected via USB. With any group of connected systems, it is possible to subvert one technology to gain access to another. An automobile will use the controller area network (CAN) bus to allow communications between the dozens of control units, including the engine control unit and onscreen display. With self-driving vehicles becoming popular, we can expect a whole new landscape.

There are potential vulnerabilities all over the place if a system isn't designed well. With so many auto manufacturers and models, the risk level only increases. Many modern vehicles also have an SD card slot, used to update GPS/maps or other automotive software. Infected SD cards could possibly be used by a person with physical access to corrupt the vehicle's GPS system with the potential for location information being sent via the operator's mobile device or from the vehicle itself. SD card slots can be disabled on some vehicles, and the settings for the on-board computer can be locked with a passcode. In the case of the CAN bus, the engine control unit is usually air-gapped. Many vehicles now use WiFi and Bluetooth, as well as proprietary technologies. When they are Bluetooth enabled, you should utilize complex Bluetooth passcodes. Also, you should consider whether WiFi and/or Bluetooth are really necessary, and if not, you can disable them within the vehicle's onboard computer settings. In an automotive CAN bus system, electronic control units (ECUs) can be the engine control unit, airbags, audio system, and so on. A modern car may have up to 70 ECUs, and each of them might have information that needs to be shared with other parts of the network.

Figure 14-3 shows CAN bus communication.

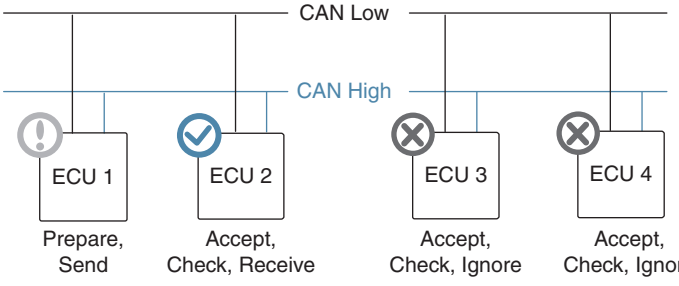


FIGURE 14-3 CAN Bus Communication

The CAN bus system enables each ECU to communicate with all other ECUs—without complex dedicated wiring. The CAN bus system consists of only two wires. An ECU can send sensor data information via the CAN bus by broadcast. The broadcasted data is accepted by all other ECUs on the CAN network, and each ECU checks the data and determines whether to receive or ignore it.

NOTE The CAN bus is simple and low cost, fully centralized, efficient, and extremely robust.

Key Topic

Aircraft

Because of the high safety and dependability requirements for aircraft systems, regulations are very strict. Software (DO-178B), hardware (DO-254), and network

(AFDX) design are all standardized. DO-178B is the primary document by which the Federal Aviation Administration (FAA) approves all commercial software-based aerospace systems. The FAA introduced AC20-171 to clarify and provide guidance to DO-178B. DO-254/ED-80 is the well-established software standard counterpart for DO-178C that contains both hardware and software, where both are critical to safe operations of an aircraft.

NOTE For more information, see www.faa.gov/aircraft/air_cert/design_approvals/air_software/media/differences_tool.pdf.

The FAA regulates all aircraft; therefore, software is a critical component of all aircraft. Recently, software glitches from a major manufacturer caused two plane crashes and the deaths of 346 people.

The autopilot is a system that serves to reduce the workload of the human pilot. Sensors produce a signal according to the direction the aircraft is flying. The control system compares this information with the desired direction. If there is a difference, the autopilot tries to correct the present situation by controlling the actuator. In the more advanced, three-axis systems, more signals are used than in the one-axis system. The control system controls both the desired orientation of the plane and the desired course of the aircraft (in direction and altitude). In that case, a small deflection from the balanced situation causes the aircraft to become unstable and uncontrollable; for example, if an attacker makes a minor trim adjustment in the balance, the pilot could be unaware of the issue, potentially drifting into another plane on descent.

There is a long list of aerospace application-embedded systems, and the list is continually getting longer. Some aircraft have embedded systems at work from tip to tail, whereas others use embedded systems in very specific ways. Embedded systems are designed in a ruggedized manner to withstand extreme conditions without any kind of loss of function. Some embedded systems perform critical functions that are responsible for passenger and operator safety. Commercial airliners provide a good example of embedded systems being used for in-flight entertainment systems, temperature control, speed control, flight management systems, flight data recorders, engine control, and phone calls, to name just a few.

The increased use of some new types of cockpit communications, including controller-pilot data link communications (CPDLC), means the cockpit is at a higher risk of hacking or interference than ever before. CPDLC enhances air traffic control (ATC) surveillance and has some intervention capability. It also plays an instrumental role in reducing mid-air collision risk, while also decreasing voice traffic on radio frequencies. Airlines have taken some steps to reduce the exposure of these systems to prying eyes, including nontransmission of the tail number or aircraft identification information, and isolation of infotainment and other critical

systems. With the help of policies and education programs for all flight personnel, they can reduce the social engineering risk.

Key Topic

Smart Meters

A smart grid is an electrical distribution system that includes a variety of operational and energy measuring devices that provide up-to-the-minute consumption, generation, and operations data to control centers. Included as part of the smart grid are smart and advanced meters. ***Smart meters*** enable control centers to view, manage, regulate, and turn down power if an issue occurs on the larger power grid and provide additional power if it is not available. Californians have seen these used for years in rolling outages. SCADA systems usually provide a back end for what smart meters are talking to; these systems collect and manage data.

In the past, attackers have been able to manipulate measured energy consumption, cause destabilization in sections of the grid, and obtain private data related to consumption usage of customers. In 2015, hackers launched an attack on the Ukraine power grid using Crash Override malware that took down an entire city for days. Smart meters are getting smarter, however. Newer versions of these embedded devices now have the capability to detect anomalies and warn of parameter changes. Even after these capabilities are widely deployed, these systems must use layered security. Utility companies have begun isolating the metering systems from their corporate networks by setting up firewalled networks, separate control system networks, and restricted access to ensure protection and privacy. Meter manufacturers need to be able to provide the right connectivity for their meters, according to strict service-level agreements (SLAs), and by ensuring security is top priority.

Figure 14-4 shows a smart meter.



FIGURE 14-4 Smart Meter

Key Topic Voice over IP (VoIP)

Voice over Internet Protocol (VoIP) is a process for sending audio signals, primarily voice, over a data network, such as the Internet. The audio is converted into a digital signal and compressed to reduce data throughput requirements. Then it's converted into packets and streamed across the network. At the receiving end, the data is decompressed and converted back to an audio signal.

There are a number of advantages to using VoIP instead of analog transmission, with the main one being that you do not need dedicated analog cables to carry the signal. In many situations, an existing data network, wired or wireless, can be used for VoIP. Routing the signals to different destinations is also much easier because it's just a matter of changing the destination address instead of physically switching the analog circuit. Because the signals are converted from analog to digital, maintaining good audio quality is easier, even in harsh industrial environments.

VoIP communications typically occur between two endpoints. Data packets can be sent directly between them and may not need an intermediary server during a conversation. In many industrial applications, the required endpoints are known and can be programmed directly into the systems. Phones are now embedded small computers that have features and capabilities like showing video, directories of users, and ticker-type data. To secure these systems, you can configure phones to connect to only specific server IP addresses, whether onsite or in the cloud. If you are hosting a call management system (CMS) on-premises, you should firewall it off from the rest of your network and allow communication only on specific UDP/TCP ports. Securing VoIP systems starts with segmenting voice traffic to dedicated VLANs and includes the following preventive measures:

- Deploy firewalls and intrusion prevention systems (IPSs) to monitor and filter authorized and unauthorized VoIP traffic, as well as track unusual voice activities on all voice traffic networks.
- Use VPNs for remote phones, especially when HTTPS/SRTP is not available.
- Enable encryption SRTP for all VoIP traffic to block eavesdroppers with sniffers and traffic monitoring tools.
- Encrypt signaling at the gateway to your provider with Session Initiation Protocol (SIP) over Transport Layer Security (TLS).
- Lock voice servers physically and logically for administration.
- Centralize administration, use domain restrictions, and enable two-factor authentication for administrative access.
- Regularly install OS updates.

- Enforce strict security policies with users.
- Limit software updates/loading on phones to administrators only.

Figure 14-5 shows a VoIP device being used through the Internet.

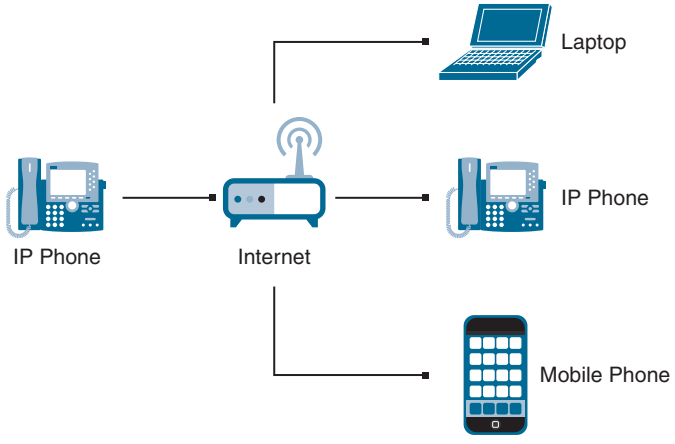


FIGURE 14-5 VoIP on a Network

Heating, Ventilation, and Air Conditioning (HVAC)

Heating, ventilation, and air conditioning (HVAC) is important for server rooms, data centers, and other technology-oriented areas of your building. Servers run hot; their CPUs can make the temperature inside the case skyrocket. This heat needs to be dissipated and exhausted outside the case. All the heat from servers and other networking equipment is enough to make your servers fail.

Key Topic

The controls for data center HVAC systems should be within the server room or network operation center and be protected by a key code. This way, only authorized IT personnel (who have access to the room) can change the temperature or humidity outside the SCADA system. These control systems can also be hooked up to the door access system or other monitoring systems to log who made changes and when. Since most servers and other IT equipment use front-to-back heat dissipation, the heat should be exhausted out behind the row.

A heating system is rarely needed in a server room, unless the organization's building is in the coldest of environments. This is due to the amount of heat that servers give off and the fact that they run 24/7. To help monitor HVAC systems and their power consumption, industrial control systems (ICSs) such as the supervisory control and data acquisition (SCADA) computer-controlled system are used. A SCADA system combines hardware monitoring devices such as PLCs (pressure gauges,

electrodes, and remote terminal units that connect to sensors) with software that is run on an administrator's (or building management employee's) workstation, allowing the administrator to monitor the HVAC system in real time. There could also be a human-machine interface that displays SCADA animations on a separate screen in a strategic place in the building.

SCADA systems can be vulnerable to viruses (such as Stuxnet) that can be used to access design files. To protect against these issues, the workstations that run the software portion of SCADA should have AV/malware software updated, and any separate physical interfaces, displays, and sensors should be secured.

Your network is only as strong as its weakest link. In the case of one department store chain, for example, an attacker gained access to the HVAC system and used it to pivot into the rest of the network by loading malware on retail point-of-sale (POS) systems to steal credit cards. In this case, approximately 40 million credit cards were collected. The obvious lesson here is that remote access should be limited and handled via VPN, secure token, and least privileged access, and monitored with intrusion detection for usage, changes, anomalies, and issues.



Drones

Embedded computers used in unmanned system applications are often characterized by their low SWaP-C (Size, Weight, Power, and Cost) profiles, small form factor (SFF), and rugged operating ranges, which are vital components for unmanned aerial vehicles (UAVs), autonomous underwater vehicles (AUVs) and remotely operated underwater vehicles (ROVs). Commercial, military ***drones***, home and hobby drones, and professional service drones all use embedded compute platforms. Drone systems rely on dedicated functions assigned to embedded computers that include video compression and streaming for drones, deep learning and neural networks for image classification and object detection in autonomous cars, and processing of sensor-acquired data such as GPS.

The FAA estimates that by 2023 there will be more than 835,000 commercial drones in the United States. As the use of drones for many commercial purposes such as aerial inspections, utility projects, real estate monitoring, and construction activities increases, more and more organizations will consider how to integrate these devices into existing networks and systems. However, these organizations must also consider how to limit the cybersecurity and privacy risks associated with the data collected by the drones.

Drones operate by using software or firmware, and drone operators use computers and mobile devices to run drone applications that control the drones. Drones communicate via wireless connections to ground stations and operators below. Hackers are already exploiting drone software and firmware vulnerabilities to take

over drones and gain access to connected systems and networks. Malware is often embedded in drone software and can compromise not only the data collected on drones, but also the systems that the drones, software, or connected devices are linked to. According to Cybersecurity and Infrastructure Security Agency (CISA, www.cisa.gov), you should ensure that devices used for the download and installation of unmanned aircraft system (UAS) software and firmware do not have access to the enterprise network and contain appropriate antivirus software. You also should properly verify and securely conduct all UAS activity on a vendor's website only and ensure file integrity using hashes. In addition, you should ensure the data link communicating with the drone supports a high encryption algorithm for securing Wi-Fi communication. Also, be sure to use the most secure encryption standards available and use complex keys that are changed regularly.

Multifunction Printers (MFP)

Key Topic

A ***multifunction printer (MFP)***/multifunction device (MFD) all-in-one (AIO) is an office machine that incorporates the functionality of multiple devices in one. The purpose is to have a smaller footprint specialized in document management/distribution/production in a large office setting. A typical MFP may act as a combination of emailer, fax, photocopier, printer, and scanner. MFPs have all of the components a computer has: hard drive, operating system, and network connection, which can be wired or wireless.

MFPs are vulnerable to attacks from your network and can be compromised. Once compromised, an MFP can be used for acquiring information about your network, which an attacker will use to penetrate mission-critical targets, such as servers. Attackers can also use the MFP for spamming, network scanning, and denial-of-service (DoS) attacks like downing a network. Threats can even come in the form of routing all scans from an MFP to an unintended user.

Securing these devices starts with password protecting them; unfortunately, many companies decline to password protect their printers due to the inconvenience it causes to users and administrators. It's important to remember that any time you send an email, fax, or remote print job through your connected office printer, data transmissions are being sent and received by the equipment. If these transmissions are not encrypted, they can be intercepted by malicious attackers, and therefore, confidential data and sensitive documentation can be stolen or compromised. Before you sell or discard old office printers, it's vital that all information and data are wiped from the internal storage. Most printers can format internal storage through the administrative control panel.

Real-Time Operating Systems (RTOS)

Key Topic

A ***real-time operating system (RTOS)*** is a software component that rapidly switches between tasks, giving the impression that multiple programs are being executed at the same time on a single processing core. In fact, the processing core can execute only one program at any one time, and the RTOS is actually rapidly switching between individual programming threads or tasks to give the impression that multiple programs are executing simultaneously.

An RTOS differs from a typical operating system in that it normally provides a hard real-time response, providing a fast, highly deterministic reaction to external events.

When switching between tasks, the RTOS has to choose the most appropriate task to load next. Several scheduling algorithms are available, including round robin, cooperative, and hybrid scheduling. However, to provide a responsive system, most RTOSs use a preemptive scheduling algorithm; one such RTOS is VxWorks.

Securing the RTOS starts with minimizing the attack surface by turning off features, services, and access that are not necessary. You should assume external systems are insecure. Do not make assumptions about what other devices the system is connected to. It is safer to assume that external devices are insecure and then approach securing them from this perspective.

Surveillance Systems

Key Topic

In the last few years, the deployment and use of video surveillance systems have increased to meet commercial needs. Uses include the industrial control industry, environmental monitoring, warehouses, facility controls, and more. These systems have contributed significantly to the reduction of crimes and the preservation of property and an increase in safety.

Far from the days of CCTV, where cameras held little to no information and were not connected to a corporate network, advances in digital video mean that connected IP cameras and associated devices on the network are at risk of being hacked and used as pivot devices. The importance of the data captured by video surveillance cameras—and what can be done with it—has led to a new breed of cyber criminals, looking for insights to steal and sell. Updated surveillance systems have many capabilities, one of which is that they can provide object and facial recognition, intrusion detection through motion, and more. Even air-gapped systems are susceptible to attacks; no system is 100 percent immune from hacks. New cyber vulnerabilities are discovered frequently, but whether they pose a critical risk depends on two factors: First is the probability that a vulnerability can be easily exploited; second is the impact that its exploitation could have on the rest of the system.

What happens when you don't take all the necessary steps but leave your data exposed to an attack? Under the EU's General Data Protection Regulation (GDPR) rules, you can be fined up to 4 percent of annual global turnover or 20 million Euros, whichever is highest.

GDPR applies also to the data captured by your security surveillance. Under the regulation, you, as network security administrator, should take all appropriate measures to provide the monitored persons with information in a brief, transparent, comprehensible, and easily accessible manner concerning the processing of their data by the camera system. The IP surveillance industry recommends you first change default usernames and passwords on all cameras and connected devices. You also should utilize the strongest encryption methods available on the devices and consider firewalling the network off from your corporate network.

System on a Chip (SoC)

Key Topic

A ***system on a chip (SoC)*** is essentially an integrated circuit or an IC that takes a single platform and integrates an entire electronic or computer system onto it. It is, exactly as its name suggests, an entire system on a single chip. The components that an SoC generally looks to incorporate within itself include a central processing unit, input and output ports, internal memory, as well as analog input and output blocks, among other things. Depending on the kind of system that has been reduced to the size of a chip, it can perform a variety of functions, including signal processing, wireless communication, artificial intelligence (AI), and more. Another way of looking at an SoC is to consider the size and capabilities of the Raspberry Pi covered earlier in this chapter, all of the functionality built into a small platform.

One of the main motivators behind the creation of systems on a chip is the fact that moving into the future, the primary goal is to reduce energy waste, save on spending costs, and reduce the space occupied by large systems. With an SoC, you can achieve all of those goals as you essentially shrink down what is normally a multichip design onto a single processor, and one that uses much less power than previous generations. SOCs have made it possible for the creation of an amazing number of portable devices that you can conveniently carry anywhere and everywhere with you without ever having to compromise on the capability and functionality of the gadgets. As such, SoCs are frequently used in systems pertaining to the Internet of Things, embedded systems, as well as cars and many more. They are used in most, if not all, portable tech such as smartphones, cameras, tablets, and other small capable wireless technologies.

Your smartphone is an example of how a system on chip works. You use your cell phone not only to make and receive calls; you also use it to browse the Internet, play games, view videos, listen to audio, take photos, and send text messages. None of this

would be possible without having multiple components combined into one, such as a graphics card, 4G/5G LTE support, wireless connections, GPS, and many other elements.

An SoC allows you to take all of these components, put them on a single chip, shrink it down to a size that can fit in the palm of your hand, and carry it around as a living component in your phone. Security on these devices has improved over the last few years and is expected to continue to innovate. Until then, many of these devices that use SoC have security concerns and capabilities, as previously discussed. To secure a SOC-based device, enable password protection and encryption, utilize complex non-patterned upper/lower/numeric/special character passwords, and require the devices to wipe themselves after a certain number of incorrect entries.

Key Topic Communication Considerations

Nearly every IoT and embedded system has some form of communication, whether it uses Bluetooth Low Energy (BLE), Wi-Fi, Zigbee, 3G, 4G LTE, or even the latest 5G. These devices are created to communicate with the outside world. This capability extends your corporate and home perimeter network to virtually an unlimited scope.

5G

Many countries have been involved in the race toward realizing 5G networks. ***5G***, or fifth-generation, is an advanced wireless network technology developed based on the 802.11ac IEEE wireless standard. It will replace its predecessor, the fourth-generation 4G LTE network, and is expected to increase the data transmission and communication by over three times. Industry experts agree that 5G will offer unlimited potential toward enhancing human life. It is considered to be a key enabler of developing more reliable and fast connections to smartphones and all other smart devices. Following are some of the key advantages of the network:

- Expanded communication capabilities, thus increasing the capacity to connect more people and devices
- A lower latency of 1 ms, which will enable users to encounter fewer lags and delays when attempting to access data through the network
- Higher rates of data, which can range between 1 and 10 Gbps, thus enabling users to download and game quickly

5G device development is still ongoing as most organizations strive to ensure their IoT devices are 5G-ready. It works alongside the existing 4G and 3G technologies to help drive an unprecedented increase in IoT innovations. It provides the platform

required to process vast data amounts to help realize a more connected and smarter world. To mention just a few applications, 5G is and will continue to connect everything, including surveillance, smart cities, and autonomous vehicles, and it can even enable a surgeon to operate on a patient in a different country in real time.

As you would expect, 5G networks also raise considerable concerns with regard to cybersecurity. 5G requires a conversion of most software-related networks; implementing future upgrades will be like installing new updates to a smartphone or computer program. This will breed numerous cyber vulnerabilities, so security professionals will have to contend with retooling tools and procedures to secure this essential network in the 21st century. From self-driving vehicles, IoT, AI, and machine learning (ML) to health care with high speeds and always-connected devices, attackers have a new playground. Securing 5G devices starts with the same familiar advice: You should disable services that are not required and configure devices with zero trust (per NIST Special Publication 800-207).

One major 5G security feature is international mobile subscriber identity (IMSI) encryption. All user data passing through 5G networks is confidentiality protected and integrity protected hop by hop. On end devices, you should ensure you use strong passwords. To help reduce your devices' exposure, be sure to disable discovery-based protocols and settings. After devices are set up and configured, disable all discovery capabilities, and disable services that are not necessary.

NarrowBand

NarrowBand-Internet of Things (NB-IoT) standards are based on low-power wide-area (LPWA) technology developed to enable a wide range of new IoT devices and services. NB-IoT significantly improves the power consumption of user devices, system capacity, and spectrum efficiency, especially in deep coverage. NB-IoT low-power consumption extends battery life to more than 10 years, which can be supported for a wide range of use cases. New physical layer signals and channels are designed to meet the demanding requirement of extended coverage (rural and deep indoors) and ultra-low device complexity. Initial costs of the NB-IoT modules are comparable to GSM/GPRS. The underlying technology is, however, much simpler, and its cost is expected to decrease rapidly as demand increases. Supported by all major mobile carriers, along with chipset and module manufacturers, NB-IoT can coexist with 2G, 3G, 4G, and 5G mobile networks. It also benefits from all the security and privacy features of mobile networks, such as support for user identity confidentiality, entity authentication, confidentiality, data integrity, and mobile equipment identification.

Baseband Radio

Baseband radio refers to the original frequency range of a transmission signal before it is modulated; it can also refer to the type of data transmission in which analog data is sent over a single non-multiplex channel. Whereas a baseband unit (BBU) processes baseband in telco systems, in a cellular phone system, a cell site has a digital signal processor (DSP) to process and forward voice signals to a mobile unit. Telecommunication networks, as we know them now, consist of four distinct parts:

- A radio access network (RAN)
- A core network
- A transport network
- An interconnect network

These networks carry three distinct traffic types, commonly referred to as planes. The control plane carries the signaling traffic, the user plane carries the user data, which is the content of communications, and the management plane carries the administrative traffic. The administrative traffic contains configuration and control commands for the RAN and core functions. Network security is critical to these planes because all three of them are prone to unique and diverse types of threats. There are serious issues in pre-5G baseband, including the ability to remotely turn on a microphone and listen in, send all calls to voicemail, remotely execute code, and crash the device. It's now trivial to set up a fake base station and have phones connect to exploit. Users should ensure they only load known signed/good applications, that their phones have the SIM chip encrypted/password protected, that they allow updates to the operating system, and that they immediately report any issues to their providers.

5G and the rapidly evolving associated technologies in the market put new security demands on telecommunication networks compared to previous mobile generations. The separation of RAN and core was critical to the evolution of 5G networks because gNBs, also known as 5G base stations, terminate the encryption of user data, except where it was previously encrypted and is outside the controls of a 5G operator's network. These capabilities allow for the enhanced security mechanisms of 5G, including the use of mutual authentication—confirming sender and receiver have an established trust and the end-to-end relationship is secure. The mandating of encryption of inter/intra-network traffic and enhanced subscriber identity protection provide increased security measures.

Subscriber Identity Module (SIM) Cards

A ***subscriber identity module (SIM)***, widely known as a SIM card, is an integrated circuit that is intended to securely store the IMSI number and its related key, which are used to identify and authenticate subscribers on mobile telephony devices like mobile phones and computers. SIM cards (see Figure 14-6) can also be used in satellite phones, smart watches, computers, or cameras. SIM cards are transferable between different mobile devices. A SIM card contains a unique serial number (ICCID), IMSI number, security authentication and ciphering information, temporary information related to the local network, a list of the services the user has access to, and two passwords—a personal identification number (PIN) for ordinary use and a personal unblocking key (PUK) for PIN unlocking. We have seen an increase in SIM jacking/swapping; that's where someone convinces an employee of a telco to move a legitimate user's IMSI number to a new SIM card for the attacker, and the attacker then uses this access to reset passwords and perform account takeovers.



FIGURE 14-6 SIM Card

Security for SIM cards includes using built-in capabilities and setting a SIM card lock; moving away from SMS 2FA codes and using an application instead; and logging in to the carrier's website and setting up a PIN and security question. In addition, certain carriers allow you to "port freeze" your number.

Zigbee

Zigbee creates flexibility for developers and end users while delivering stellar interoperability. Created on the IEEE's 802.15.4 standard, using the 2.4-GHz band and a self-healing true mesh network, the technology defined by the Zigbee specification is intended to be simpler and less expensive than other wireless personal area networks (WPANs), such as Bluetooth or more general wireless networking such as WiFi. Zigbee is commonly implemented in wireless light switches, home energy monitors, traffic management systems, and other industrial equipment that requires

short-range low-rate wireless data transfer. Its low power consumption limits transmission distances to 10–100 meters line-of-sight, depending on power output and environmental characteristics. Zigbee devices can transmit data over long distances by passing data through a mesh network of intermediate devices to reach more distant ones. Zigbee is typically used in low data rate applications that require long battery life and secure networking. Zigbee networks are secured by 128-bit symmetric encryption keys. And Zigbee has a defined rate of 250 kbps, best suited for intermittent data transmissions from a sensor or input device.



Embedded System Constraints

Many embedded systems have substantial design constraints over desktop computing applications. The combination of cost pressure, long lifecycle, real-time requirements, reliability requirements, and design culture nonstandards can make it difficult to be successful applying traditional computer design methodologies and tools to embedded applications. Embedded systems in many cases must be optimized for lifecycle and business-driven factors rather than for security and maximum computing throughput. Knowing the strengths and weaknesses of current approaches can set expectations appropriately, identify risk areas, and suggest ways in which to build and meet industrial needs.

Power

Embedded systems (ESs) are typically designed for a specific purpose or task, and are optimized to meet certain constraints, performance, timing, power, and cost. Power consumption is one of the main design constraints for embedded electronic systems, such as wireless sensors, compute systems, and controller devices. Embedded systems also need to be concerned about the power dissipation ion when executing software instructions; a designer must balance the power consumption with features and capabilities. Stronger encryption levels will execute more CPU cycles, drawing more power. Extensive software instructions also consume more power. An attacker can simply cause a loop instruction somewhere on the system and cause the system to consume all power and shut down or crash. When deploying these devices, you should consider securing power sources and ensuring adequate power during the highest utilization periods.

Compute

Many embedded computers are physically located within some larger component. Therefore, their form factor may be dictated by form factors existing in pre-electronic versions or having to fit into interstices among mechanical components. In transportation and portable systems, weight may be critical for fuel economy or human

endurance. For example, a mission-critical system has much more stringent size and weight requirements than others because of its use in a flight vehicle.

Network

A computer is often embedded to interact with the environment for the purpose of monitoring and controlling external devices. To do this, analog inputs and outputs must be transformed to and from digital signal levels. Also note that significant current loads may need to be switched to operate motors and actuators, all these requirements can lead to the need for a larger computer. In some systems, smart sensors and actuators that contain their own analog interfaces, power switches, and small CPUs are often used to offload interface hardware from the central embedded computer. This brings the additional advantage of reducing the amount of system wiring and number of connector contacts by employing an embedded network rather than a bundle of analog wires.

Crypto

One of the first steps in building a secure embedded system is to see whether cryptography is actually needed. Cryptography is an important part of securing applications, but it is not a security panacea, nor is it always necessary for building a secure system. The constraints on cryptography in building secure embedded systems hardware and software have to do with the amount of code required to implement a secure algorithm and the processing power required to crunch the number. Also public key infrastructure (PKI) and certificates require space to store the certificates and require CPU cycles to perform operations. The limits of the platform need to be tested to ensure full compliance.

Inability to Patch

Patches for post-release bugs in OS software have shown that nearly 25 percent of all fixes are incorrect and directly impact the end user. And if that's not bad enough, 40 percent of these faulty fixes resulted in crashes, hangs, data corruption, or additional security problems. Patches don't always solve the security issues they were designed to address. The Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) saw a 60 percent failure rate in patches fixing the reported vulnerability in control system products. Even good security patches can cause issues because most patches require shutting down and restarting the manufacturing process. Some can also break or remove functionality previously relied on by the control system.

Authentication

Similar to cryptology implementations, which start with limits of embedded system CPU, memory, and flash/storage components, the impact of certain authentication types must be considered. Authentication, digital signatures, and encryption are generally based on a set of keys and algorithms for transforming digital data, called clear text, into an encrypted form and back again. Digital signatures are based on the encryption of a checksum of the data being signed. Secure authentication is accomplished using digital signatures, and all of these functions take a toll on memory and CPU.

Range

Embedded systems are typically constrained in terms of three system performance criteria: space, time, and energy. Performance requirements are directly translated into constraints imposed on the system's resources, such as code size, execution time, and energy consumption. These resource constraints conflict with each other in a complex manner, making it difficult for a system developer to apply a well-defined design methodology in developing a real-time embedded system. The smaller the device, the more direct relationship to the size of the transmitter, thereby dictating the frequency and distance range the device supports. Obviously, the smaller the transmitter, the shorter the distance.

Cost

The emergence of multimedia and wireless applications as growth leaders has created an increased demand for embedded systems. Applications range from digital cellular phones to high-performance avionics systems. Along with the increased market share of these products comes an increase in system complexity and cost. These systems consist of programmable-multiprocessor arrays to meet performance requirements. Time-to-market and lifecycle costs are key factors in the success of these products in the competitive electronics marketplace. These costs have a dominant influence on the design of embedded microelectronic systems. In addition, these systems must meet rigid form factors, such as size, power, and weight constraints, which further complicate the design task. For designers of high-end embedded microsystems or large-volume consumer products, rapidly prototyping cost-effective implementations that meet stringent performance, functional, timing, and physical requirements is a formidable challenge.

Implied Trust

A trusted computing platform is a computing infrastructure that provides a variety of hardware-supported and software security functions. Although *trust*, *trusted*, and

trustworthiness are never defined, it is hoped that trusted computing platforms and resulting improvements in the security of computing infrastructure and applications will enable trustworthy applications and systems. Implied trust is where a specific section of code or a component depends on the previous code or component for its security, thereby making it reliant on that code to have appropriately performed the security measures required to make the entire system safe. When it comes to cybersecurity, implied trust should be used sparingly as it relates to embedded systems. There are constraints that limit what full-fledged trust can be implemented, let alone what can be operationally expedient; therefore, certain aspects of the design rely on an implied trust, even if that is "assuming" the wireless network will be secure so that there is no need to protect plaintext protocols.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 14-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 14-2 Key Topics for Chapter 14

Key Topic Element	Description	Page Number
List	Embedded Systems	339
Section	Supervisory Control and Data Acquisition (SCADA)/Industrial Control Systems (ICS)	341
Paragraph and list	Internet of Things (IoT)	344
Section	Medical Systems	347
Section	Vehicles	347
Section	Aircraft	348
Section	Smart Meters	350
Section	Voice over IP (VoIP)	351
Paragraph	Heating, Ventilation, and Air Conditioning (HVAC)	352
Section	Drones	353

Key Topic Element	Description	Page Number
Section	Multifunction Printers (MFP)	354
Section	Real-Time Operating Systems (RTOS)	355
Section	Surveillance Systems	355
Section	System on a Chip (SoC)	356
Section	Communication Considerations	357
Section	Embedded System Constraints	361

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

embedded systems, Raspberry Pi, Arduino, Field-Programmable Gate Array (FPGA), supervisory control and data acquisition (SCADA), industrial control system (ICS), logistics, Internet of Things (IoT), sensors, smart devices, wearables, facility automation, smart meters, Voice over Internet Protocol (VoIP), heating, ventilation, and air conditioning (HVAC), drones, multifunction printer (MFP), real-time operating system (RTOS), system on a chip (SoC), 5G, NarrowBand-Internet of Things (NB-IoT), baseband radio, subscriber identity module (SIM), Zigbee

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is an Arduino device?
2. What is the purpose of an FPGA?
3. What systems do the SCADA/ICS control systems actually control in the manufacturing process?
4. Look around your house. What IoT-related devices do you own? What is the number one problem with IoT devices being developed and sold?
5. Today's vehicles are mostly computerized. Which protocol/system do most vehicles use for ECU to communicate between themselves?



This chapter covers the following topics related to Objective 2.7 (Explain the importance of physical security controls) of the CompTIA Security+ SY0-601 certification exam:

- Bollards/barricades
- Access control vestibules
- Badges
- Alarms
- Signage
- Cameras
  - Motion recognition
  - Object detection
- Closed-circuit television (CCTV)
- Industrial camouflage
- Personnel
  - Guards
  - Robot sentries
  - Reception
  - Two-person integrity/control
- Locks
  - Biometrics
  - Electronic locks
  - Physical locks
  - Cable locks
- USB data blocker
- Lighting
- Fencing
- Fire suppression
- Sensors
  - Motion detection
  - Noise detection
  - Proximity reader
  - Moisture detection
  - Cards
  - Temperature
- Drones
- Visitor logs
- Faraday cages
- Air gap
- Screened subnet (previously known as demilitarized zone)
- Protected cable distribution
- Secure areas
  - Air gap
  - Vault
  - Safe
  - Hot aisle
  - Cold aisle
- Secure data destruction
  - Burning
  - Shredding
  - Pulping
  - Pulverizing
  - Degaussing
  - Third-party solutions

Understanding the Importance of Physical Security Controls

In a facility, nothing is more important than keeping your workforce safe from any and all potential injuries. When heavy machinery operates in close proximity to your employees, damage prevention products should always be leveraged to offer total warehouse protection.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 15-1 lists some of the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 15-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Access Control Vestibules	1–2
Cameras	3
Personnel	4–5
Locks	6–7
Sensors	8
Secure Areas	9
Secure Data Destruction	10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following access controls allows security guards to view visitors prior to access engagement and entry?
   a. Access control vestibule
   b. Bollard
   c. Barricade
   d. Guard station

2. Which of the following helps control access in and out of a facility?
   a. Fencing and guard gates
   b. Bollards and barricades
   c. Access control vestibules
   d. Signage and dogs

3. Which motion detection method uses infrared or laser technology as its primary method?
   a. Radar detection
   b. Optical detection
   c. Object detection
   d. Proximity detection

4. What sort of guard can patrol an area 24/7/365 with very little downtime?
   a. Patrol cars
   b. Reception desk
   c. Armed guards
   d. Robot sentries

5. The reception desk plays a critical role in physical security controls of most companies. How do receptionists accomplish such a critical role?
   a. They greet visitors with a smile and ask their names.
   b. They place visitors in a vestibule until they provide ID.

c. They provide visitors with a tracking beacon.
d. They verify the identity of visitors, contact the employees being visited, and hold guests until they are picked up.

6. Which type of lock uses a fingerprint to unlock a device?
   a. Cable lock
   b. Key lock
   c. Multifunction lock
   d. Biometric lock

7. Which lock type requires a skeleton key?
   a. Warded/lever lock
   b. Skeleton door
   c. Master lock
   d. Cyclone dutch

8. Which type of sensor can read when something is nearby?
   a. Moisture
   b. Noise
   c. Proximity
   d. Temperature

9. Which place can you safely and securely have people working inside?
   a. Safe
   b. Crate
   c. Vault
   d. None of these answers are correct.

10. Which secure data destruction method is best for hard drives and other forms of magnetic media?
    a. Burning
    b. Degaussing
    c. Rivets
    d. Pulverizing

Foundation Topics

Bollards/Barricades

Key Topic

A ***bollard*** is a standalone post, typically steel, short, and sturdy and anchored in a hard surface such as concrete (see Figure 15-1). This low-profile, post-shaped deterrent is built with the purpose of blocking vehicle movement from certain directions, while allowing for full pedestrian movement. Bollards are designed to withstand high impacts and deflect potential blows away from the object or area they are in place to protect.



FIGURE 15-1 Bollards

Various bollard types offer different levels of protections. There is a rating system for the standards by which these bollards and barriers are judged. The standards are based on perpendicular impact at a certain speed (K), with the highest rating indicating the strongest protection, such as a vehicle ramming this bollard/barrier at a higher speed.

Antiram bollards are security bollards that prevent vehicles from crashing through them. There are different levels of antiram bollards. K4, K8, and K12 ratings offer various levels of protections against different size vehicles.

Safety bollards give a warehouse or distribution center a path/road definition, diverting traffic and intrusions. If traffic flow matters in your warehouse, these bollard types matter; they also are used inside warehouses to protect people, goods, and equipment.

Architectural and landscaping bollards offer a similar function. They may differ in that they are more focused on aesthetics than function, though function is still achieved.

Barricades or barriers are just like guardrails (see Figure 15-2). Typically used as a material-handling solution, they offer more linear protection than bollards. Barriers prevent both people and vehicle movement across a large area. They are typically constructed of steel or concrete and are designed with the goal of protection.



FIGURE 15-2 Barricade

Barriers are used to protect wide stretches of racking and walls in a facility. They are mounted to the floor, with a small amount of space between the barriers and protected assets in case of impact. They also are used at the ends of racks and to protect utilities, around machines and conveyors, and anywhere you want to designate and segregate a space. Guardrail barriers are typically painted a highly visible yellow color and can be used to define the working areas in a facility.

Key Topic Access Control Vestibules

Security vestibules are an excellent access control addition to any facility protection plan. These are purpose-built entries with panels built from prefabricated composite or metal and are used as a way for companies to control the traffic flow in their facilities, providing security with a clear view of those entering. The vestibule is configured to ensure one of the doors to the entrance is always closed, which also reduces the hot/cold air that escaping from the building. These vestibules are especially useful for companies or buildings with high traffic if a door is opened frequently. In much the same way, vestibules can boost a facility's security, as illustrated in Figure 15-3.

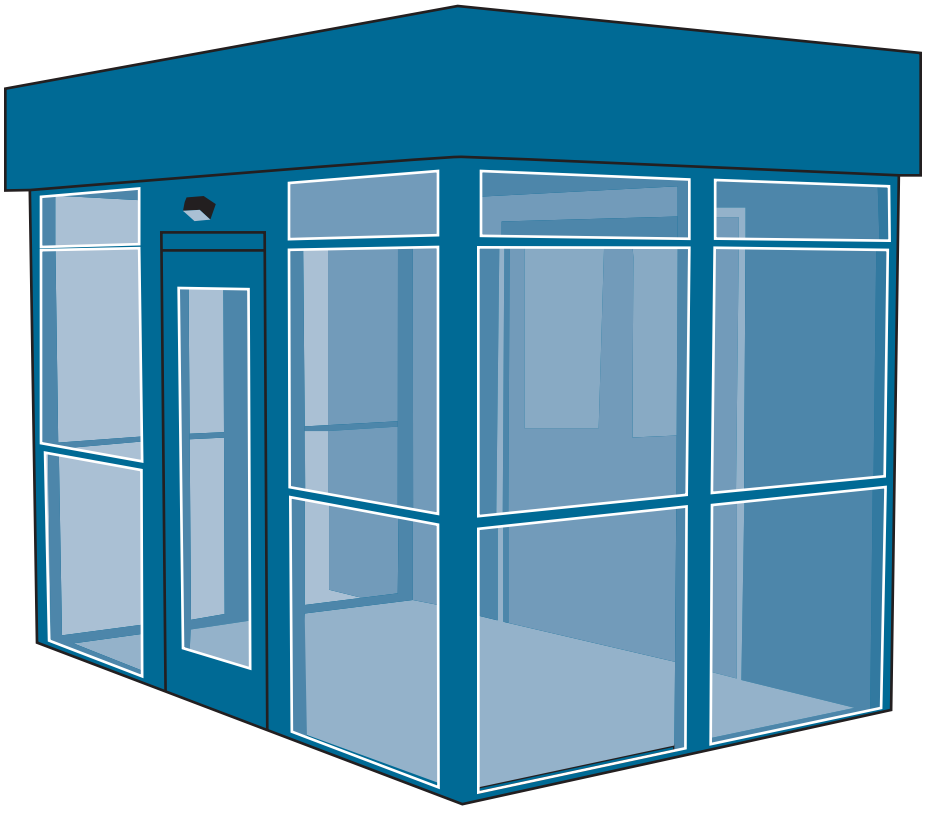


FIGURE 15-3 Access Control Vestibule

Access control vestibules provide a buffer zone to the entrance of a facility, as well as an additional layer of physical security and access control. Because they use interlocking doors, a trap is created; both sets of doors cannot be opened at a single time.

With the vestibule/trap entry being monitored and controlled at all times, your security team can ensure the person attempting to enter the building has the clearance to do so. Security vestibules can also incorporate additional security measures into the design, such as PIN pads or badge scanners, depending on your facility's needs.

Key Topic

Badges

An access badge is a credential used to gain entry to an area having automated access control entry points. Entry points may be doors, turnstiles, parking gates, or other barriers. Access badges use various technologies to identify the holder of the badge to an access control system. The most common technologies are proximity readers, barcodes, smart cards, and various biometric devices. Magnetic strip cards are fading from use. The access badge contains a number that is read by a card reader. This number is usually called the facility code, and it is sent to an access control system—a computer system that makes access control decisions based on information about the credential. If the credential is included in an access control list, the access control system unlocks the controlled access point. The transaction is stored in the system for later retrieval; reports can be generated showing the date/time the card was used to enter the controlled access point.

A reader radiates a one- to five-inch electrical field around itself. Cards use a simple LC circuit. When a card is presented to the reader, the reader's electrical field excites a coil in the card. The coil charges a capacitor and in turn powers an integrated circuit. The integrated circuit outputs the card number to the coil, which transmits it to the reader. The transmission of the card number happens in the clear; it is not encrypted. With basic understanding of radio technology and of card formats, proximity cards can be hacked.

A smart card or integrated circuit (IC) card is a physical electronic authorization device that is used to control access to a resource. It is typically a credit card–sized plastic card with an embedded IC chip, as illustrated in Figure 15-4. Many smart cards include a pattern of metal contacts to electrically connect to the internal chip. Others are contactless, and some are both. Smart cards can provide personal identification, authentication, data storage, and application processing. Applications include identification, finances, mobile phones (SIM), public transit, computer security, schools, and health care. Modern credit cards have this IC built in for stronger protection. Smart cards may provide strong security authentication for single sign-on (SSO) within organizations.



FIGURE 15-4 Smart Card

Alarms

Key Topic

Alarms are the most basic definition of any security system. They are literally a means or method by which something is secured through a system of interworking components and devices. In this instance, we're talking about facility security systems, which are networks of integrated electronic devices working together with a central control panel to protect against burglars and other potential intruders.

A typical security alarm system includes

- A control panel, which is the primary controller of a security system
- Door and window sensors
- Motion sensors, both interior and exterior
- Wired or wireless security cameras
- A high-decibel siren to draw attention
- A premises sign and window stickers

Alarms provide greater protection when they are tied into the badge and control systems. If unauthorized cards or users try to gain access to a room or facility that they are not authorized to be in, the alarm can be silent and warn the security desk or can be audible and warn these persons that their presence has been noticed.

Signage

Key Topic

Security sign placement provides an extra layer of security in addition to an active security system. When it comes to deciding what type of security measures you want to take, the first step is assessing your property type and risk of loss. If you are protecting a private property, even a vague sign referencing a security system is

usually enough to deter break-ins and burglaries. After all, thieves are looking to get in and out as quickly as possible. Having signs around a facility indicating surveillance can be an effective way to deter impulsive burglars. Strategically placing signs near back entrances and easy-access windows helps improve overall safety. Thieves will check for unlocked doors and windows first, trying to avoid breaking glass or making too much noise. Seeing a sticker that indicates an alarm will sound off once opened instills fear and is likely to send the perpetrators away. Even a "beware of dog" sign has been proven to deter potential threats, regardless of whether a dog is actually present. Beyond security signage at a facility, other signage provides direction and guidance for staff and visitors, ensuring expectations are clear and identifying the repercussions for failure to abide by those signs, as illustrated in Figure 15-5.



FIGURE 15-5 Signage

Key Topic

Cameras

As mentioned in Chapter 13, "Implementing Cybersecurity Resilience," security cameras can triple your monitoring area from a single point to many. In the physical security world, access control is used to restrict access by unauthorized personnel to a property, building, or room. IP video cameras are a type of digital video camera used for surveillance; they differ from analog closed-circuit television (CCTV) cameras by being able to send and receive data over a computer network or via the Internet.

There are two types of IP cameras: centralized and decentralized. Centralized cameras require a central network video recorder (NVR), whereas decentralized cameras

have built-in recording functionality and store video locally on flash memory, network attached storage (NAS), or hard disk drives.

Other components that are features of security cameras include

- ***Motion recognition*:** Newer cameras and software can recognize that something moved or changed position relative to its surroundings or the surroundings of the object. The two principal methods by which motion can be electronically identified are optical detection and acoustic detection. Infrared or laser technology is primarily used for optical detection. Optical detectors convert incoming optical energy into electrical signals. The two main types of optical detectors are photon detectors and thermal detectors. Photon detectors produce one electron for each incoming photon of optical energy. The electron is then detected by the electronic circuitry. Motion detection devices, such as passive infrared (PIR) motion detectors, have a sensor that detects a disturbance in the infrared spectrum. Once detected, a signal can activate an alarm or a camera that can capture an image or video of the motion and start recording the event.
- ***Object detection*:** Certain cameras and their systems can evaluate and detect objects and the world around them. By constantly scanning and evaluating a specific area, a system can be trained to classify objects and detect when a change has taken place. A lot of newer facial recognition systems now have this capability, and although this newer field is still being developed, there are a lot of applications. When evaluating an area, You Only Look Once (YOLO) is the fastest method and therefore most-used real-time object detection system. Basically, it applies a single process to the full image, dividing it into regions. Then the network predicts bounding boxes and probabilities for each region; this allows the system to make comparisons with previous items (trained) on the system.

Key Topic Closed-Circuit Television (CCTV)

Closed-circuit television (CCTV) is mainly an older analog technology, where you build out a coaxial cable network and connect cameras to the cable. All cables lead back to a distribution system and then to a DVR and security monitors. One of the issues with CCTV is that it requires you to build a completely separate network, one that may be alongside your IP (data) network. In industrial plants, CCTV equipment may be used to observe parts of a process from a central control room—for example, when the environment is not suitable for humans. CCTV systems may operate continuously or only as required to monitor a particular event.

Newer CCTV can use digital video recorders (DVRs), which provide recording for hours to many years, with a variety of quality and performance options and

extra features like motion detection. Most CCTV systems are being replaced by IP camera networks although they are still being called CCTV; the concept behind CCTV is that it is a surveillance tool. Cameras, whether they are real or dummy, provide a deterrent from attackers because they are fearful of being recorded and later prosecuted. Placing cameras at egress and ingress spots around facilities ensures security guards are able to attend to issues that may be missed by the cameras, and they provide backup and evidence for solving crimes.

Key Topic

Industrial Camouflage

The goal of ***industrial camouflage*** is to deceive the view of the viewers, in an attempt to provide them with either a blended image where the surrounding scenery and the camouflaged structure appear as one, or to have them believe the structure is something else entirely. For example, roofs are planted with grass or even with small shrubs that match their surroundings; placing trees strategically provides even better concealment and natural irregularity of light and shadows.

During World War II, after the attack on Pearl Harbor, the Lockheed plant put in place an extremely elaborate visual deception to convince planes at 5,000 feet or greater to see fields and houses, by painting parking lots various colors of green and by scattering fake house roofs around the top of the plant, along with camouflage chicken wire and painted feathers. In modern times, military forces wear clothing and machinery with camouflage patterns. Likewise, data centers surround their premises with trees, bushes, and other vegetation and implement low-profile security measures around the perimeter to blend in with the surroundings.

Personnel

The people a company hires are its number one asset; they are also included among the company's other lines of defense. Hiring the right people goes beyond assessing their capabilities, expertise, and ingenuity. Being loyal to a company may be a thing of the past, but companies must build programs to educate employees about protecting their jobs by protecting the company's intellectual property. Hiring guards, sentries, and reception desk personnel who put security first is a must; they all act together to offer the following security benefits:

Key Topic

- ***Guards*:** Among a company's first lines of defense/offense are security guards. Guards who are deployed at ingress/egress points can reduce tailgating, verify the identity of scanned ID cards, and ensure that property does not leave the premises. Guards patrol the internal and external building space and ensure unauthorized persons are apprehended or removed.

- ***Robot sentries*:** Robot sentries act as 24/7/365 guards, continuously monitoring and alerting on differentials. They report anything out of the ordinary to the security desk/office, where the security team can deploy additional personnel to address any potential breach or other issues. Robots help extend the monitoring capabilities of security guards by having an always-watchful eye.
- ***Reception*:** The reception desk places a friendly face in front of visitors, essentially a buffer zone keeping strangers from accessing anyone in the company. Receptionists vet the people attempting to gain access to the facility or particular employees. They do this in a number of ways—first by checking IDs, next by verifying with the employees, and also by providing visitor badges. These visitor badges alert employees as to where these persons are or are not supposed to be. Receptionists contact the appropriate person to verify visitors are indeed in the facility for a meeting, for example, and keep the visitors at the registration desk until they are picked up by the employee meeting with them. This line of defense prevents unknown people, including potential attackers, from roaming the halls and accessing employees' desks, equipment, and other critical parts of the building.
- ***Two-person integrity control*:** Integrity controls usually include an escort who is assigned to a particular person doing work in a high-risk area, such as a data center, where a single person can cause massive amounts of damage in a few minutes. One of the two people is there as an observer (or escort); this person monitors the one performing the work and ensures that person is performing work exactly as described in the change request. The monitor can also question any variance and can stop the implementer if he or she feels that person is not following guidelines and will report any variance or unusual or suspicious activity immediately to supervisors.

Locks

Locking all doors, windows, and ingress/egress points of a facility is not only good advice for business but also for homes. Open doors, windows, and gates are an invitation for mischief. Security computer cables can lock laptops, projectors, and other assets so that they cannot be removed from the facility. By locking portable and expensive assets, you make them unmovable, thereby protecting them from being taken and protecting the data on those devices. The following options are available for locks:



- ***Biometrics*:** Covered in earlier chapters in greater detail, biometrics provide a unique way of making sure people are who they say they are by monitoring/matching human characteristics such as a fingerprint, retina, or voice. Biometric controls and locks add to the security-layered approach to physical security controls.

- ***Electronic locks*:** These locks usually have a magnetic strike plate and are energized to keep the magnet engaged until authorization is given to disengage, thereby releasing the magnetic hold. Most newer electronic locks—even those for homes—now require a PIN code, a fingerprint, a retina scan, or an ID badge. Electronic locks can be configured to report when opened or closed and who accessed them and when. They can also be controlled by a security desk so that access is denied until the appropriate identity is confirmed.
- ***Physical locks*:** These locks have been around for hundreds of years; they can have digital codes, tumblers, keys, or sliding dials with numbers. Locks have levers and interior notches to correspond with the ward (notch); these are called lever or warded locks. They protect the front door of many businesses; they even protect your home. These locks vary in style, type, and capability. Many of them are trivial to pick, so anyone with basic lock-picking skills can easily pick most of them. Newer, anti-lock-picking locks with complex tumblers and drop-down pins ensure these locks are nearly invincible.
- ***Cable locks*:** These locks are typically deployed to attach secure items to stable, heavy equipment in a room; that is, something that is not easily removed. You could use a computer cable lock that connects a laptop computer to a desk or a projector to its cart or stand. Cable locks have been used to lock bikes to bike stands or fences for a very long time. Although these devices don't stop determined thieves, they do slow them down and could allow enough time for another monitoring/deterrent system to kick in.



USB Data Blockers

Anywhere there is intellectual property or data classified at a high level, you will find that computers have their USB ports disabled, or they are enabled to support only a keyboard and mouse. In the past we have seen where newer attack tools can mimic a mouse or keyboard; fortunately, many companies and government agencies are utilizing security device lockdown software called data blockers that can detect the emulation USB device that has been placed in the device and can promptly disable usage, as well as report the exploit to the security operations center (SOC). USB data blockers ensure people, employees, visitors, and attackers are unable to steal data with a USB thumb drive from a corporate computer by disabling the USB function.

Another type of ***USB data blocker*** is a USB charge-only device. These are not the same as the other blockers. USB charge-only cables and devices are typically used to protect a user's phone or personal items from being attacked by a back-end system that is capable of reading the device input and running an attack to exploit the device to steal data. These USB data blockers or charge-only cable/devices have

several pins removed so that data is unable to traverse the USB cable. You may see these free charging stations at airports, hotels, and office lobbies. If you have the appropriate cable or device (charge only), then you can feel safe to use it; otherwise, it is always best to use a normal power plug. These USB data blockers have also been called USB condoms.

In a typical attack that has been successfully executed many times, an attacker drops backdoored USB drives in a common area, parking lot, or break room. When someone picks up the USB drive and inserts it into a system, it automatically launches and installs malware.

Lighting

Key Topic

Lighting inside and outside a building can provide various levels of protection and safety for employees and the facility. Knowing there is lighting everywhere can also serve as a deterrent for potential attackers because it's possible for cameras to pick up any movement. Planning is required to design and deploy lighting appropriately to ensure there are no dark spots and to provide safety for employees and the facility.

Building lighting and autosensing light switches can provide safety and visibility as well as reduce power requirements. Newer lighting is direct current (DC) based; because lights are now mostly LED based and do not require 110V, they can operate on 5V. As a result, entire DC power distribution systems have been developed.

Employee and visitor parking lot lighting is vital to the safety of your employees and guests. You must ensure that lights in these areas overlap sufficiently to avoid dark spots or shadows. Having higher-power lighting at ingress and egress spots that are on separate controls and can be used during attacks or emergencies provides highly visible paths to certain spaces.

Fencing

Key Topic

Physical access controls consist of the systems and techniques used to restrict access to a security perimeter and provide boundary protection. One of the major defenses includes fencing and fence-monitoring systems. Fencing is usually the first line of defense at the perimeter of a property. Therefore, deploying the right fence for the right level of protection is important. Today's fences include crash barrier bollards and control lines to ensure that a rushing vehicle cannot pass. Fence-line trembler devices can alert the security desk to motion on the fence. When this approach is combined with other visible methods, such as CCTV, the threat is visible immediately and the impact of the threat is reduced.

Pedestrian barriers should include controlled entry and exit points, locked gates, turnstiles, and electronic access control systems. Fencing can be as simple as wood

or chain link or even brick, steel, or wrought iron. Remember, the stronger the fence material, the less likely it will be broken or traversed by criminals. Some safety fences are even made of vinyl and plastics to dissuade people from entering areas that are dangerous or under construction, but you need to ensure signage is appropriate.

Key Topic

Fire Suppression

A ***fire suppression*** system is an engineered group of units that are built to extinguish fires through the application of a substance. Most commonly, a fire suppression system has built-in components that detect fires at the beginning stages through heat, smoke, and other warning signals. These are attached to an alarm system that will alert you when the fire has been detected and initiate steps for action to further suppress the fire. Most fire suppression systems automatically release a substance to extinguish the fire after the detection and/or alert. The most common clean agents used to suppress fires are FM200 and Novec 1230; these chemicals do not destroy components like water does.

Most fire suppression systems have a manual application release. It is important to secure this button behind an easy-to-reach-and-activate plastic flip-up wall-mounted box. This setup ensures that the person pressing the button is doing so consciously and not by accidently bumping into it.

When fire suppression systems are triggered, they are usually tied to power cut-off processes, which turn down power at the same time. This reduces chances of the fire continuing to burn while sourced by a device with power, and of someone inadvertently getting electrocuted. Class C fires (energized electrical equipment, electrical fires, and burning wires) are put out using extinguishers based on carbon dioxide. For Class A fires (trash, wood, and paper), water decreases the fire's temperature and extinguishes flames. For Class B fires (flammable liquids, gases, and greases), foam extinguishes flames. Class D fires involve combustible metals such as magnesium, titanium, and sodium. The two types of extinguishing agents for Class D fires are sodium chloride and a copper-based dry powder.

Sensors

The deployment of sensors is a cost-effective way to enhance and extend monitored areas. They are capable of sensing and monitoring for everything from window and door openings by attackers, to water and gas leaks, to motion detection and much more. Every organization deploys them in some manner. When used in conjunction with other deployed security measures, they can add an extra level of security in a reasonable manner. Many sensors enhance protections by providing notification/alerting for components; they include

Key Topic

- **Motion detection:** Motion detection sensors are a linchpin of any security system because they detect when someone is in a location that is not expecting to see motion. These detectors use passive infrared to detect motion and then send an alert to the alarm system or security desk.
- **Noise detection:** Noise detection or sound sensors can be anything from glass break sensors to room monitor sensors. If sound is detected where it is not expected, that sound is converted into an electronic signal and then sent to an alarm or alerting system to notify the appropriate security person.
- **Proximity reader:** A proximity reader or prox card reader, typically an RFID reader, reads a prox card by placing it near (within proximity of) the reader. The reader sends energy in the form of a field to the card, which powers up the card, allowing the proximity reader to read the information stored on the prox card. Prox cards are used as part of an access control system, and when tied to other verification systems like video, proximity readers can help provide a layered approach to providing physical access to facilities.
- **Moisture detection:** Moisture detection systems look at water and humidity under the floor in data centers or in other critical spaces to provide early warning of a leak, flood, or some other release that could cause substantial damage to life and property.
- **Cards:** As previously mentioned, badges, ID cards, and access cards are all part of access control systems, making sure that the appropriate person has access to the right place, and unauthorized persons are immediately identified and removed.
- **Temperature:** Temperature sensors are placed in data centers, racks, and throughout buildings. In data centers and racks, they ensure equipment doesn't overheat and cause catastrophic failure of critical assets. In some parts of the building, they work with the HVAC system to ensure it is not too cold or too hot for employees.

Key Topic

Drones

Drone security starts with physical security controls, controlling access to the Drone control systems with protected door access. Authentication methods need to ensure proper access to not only the systems that control the drones but also to the campus, building, and room where the controls are housed. Organizations can put into place drone detection, monitoring, flight restrictions, and jamming to ensure only authorized drones are allowed on or near the facility. You can also set up geofencing on company drones to ensure that they don't drift outside restricted space. To operate a drone, you are required to register with the Federal Aviation Administration (FAA)

and abide by all laws and public safety rules. Drones can be programmed, automated, and configured to patrol the perimeter of your facility, freeing up security guards and giving you a much wider, broader, and higher field of view than a single security guard could ever cover.

Key Topic Visitor Logs

A visitor logbook has many benefits that your organization can use to its advantage. Although some might consider a log "just another inconvenience," when used properly, a visitor log can help protect both your organization and your guests and provide other valuable information. Whether they are paper or digital visitor logs is a matter of preference, although one has obvious advantages over the other. Numerous visitors pass through the average workplace on any given day. Customers, clients, delivery personnel, family, job candidates, and more—they all stop in for various reasons and may stay for only a few minutes or for most of a day. The primary benefits include the following:

- **Security:** These logs provide access control and heightened overall security that maybe perceived or real.
- **Emergency preparedness:** The logs enable you to know who is onsite and ensure they are counted in assembly areas.
- **Analytics:** Data of comings and goings can be valuable for a number of reasons.

Visitor logs are part of an effective visitor management system, and visitor management often starts with access control. To be truly effective, organizations should control the visitor registration process by requiring all visitors to enter through one specific entrance. This restriction further prevents unauthorized access to the building by visitors who have not been vetted in any way. One entrance is also more user friendly; guests know exactly where they are supposed to be and so do security guards and employees.

Key Topic Faraday Cages

A ***Faraday cage*** is an enclosure used to block electromagnetic fields. A cage or shield can be formed by a continuous covering of conductive material, or in the case of a Faraday cage, by a mesh of such materials.

A Faraday cage operates because an external electrical field causes the electric charges within the cage's conducting material to be distributed so that they cancel the field's effect in the cage's interior. The Faraday phenomenon is used to protect sensitive electronic equipment such as RF receivers from external radio frequency

interference (RFI) during testing or alignment of the device. They are also used to protect people and equipment against actual electric currents such as electrostatic discharges and lightning strikes, since the enclosing cage conducts current around the outside of the enclosed space and none passes through the interior. Faraday cages can shield the interior from external electromagnetic radiation if the conductor is thick enough and any holes are significantly smaller than the wavelength of the radiation. Companies have found placing expensive critical equipment in racks that are Faraday purposed provide an extra layer of protection from free-range electromagnetic discharge or radiation. Faraday cages/rooms are spaces that are completely enclosed by one or more layers of a metal mesh, perforated sheet metal, or fence-like material. The metal layers are grounded to dissipate any electric currents generated from external or internal electromagnetic fields, and thus they block a large amount of the electromagnetic interference.

Key Topic Air Gap

Air gap, or "gapping," refers to the concept that there is a gap of air or no connection between the computer and other networks. Because the computer isn't directly connected to the network, it can't be attacked through the network. So, to compromise this type of computer, hackers have to "cross the air gap," which means they need to be physically sitting down in front of the computer. An air-gapped computer system has no physical connection to other computers, networks, or unsecured systems. As security administrator, you should consider disabling ports and enforce a policy of not using USB thumb drives. This way, you can reduce the likelihood of a virus or data theft.

Key Topic Screened Subnet (Previously Known as Demilitarized Zone [DMZ])

In computer security, a ***screened subnet*** or demilitarized zone (DMZ), often referred to as a perimeter network, is a physical or logical subnetwork that contains and exposes an organization's external-facing services to an untrusted, usually larger, network such as the Internet. The purpose of a screened subnet is to add an additional layer of security to an organization's local area network (LAN). An external network node can access only what is exposed in the screened subnet, while the rest of the organization's network is protected by firewalls. The screened subnet functions as a small, isolated network positioned between the Internet and the private network.



Protected Cable Distribution

The purpose of a ***protected cable distribution system***, also called a protective distribution system (PDS), is to deter, detect, and/or make difficult physical access to the communication lines carrying data and/or voice communications. In data centers, this can mean specifically hardened rooms that are separate and apart from other rooms, and access control measures are in place to ensure appropriate monitoring and access.

In a hardened protected cable distribution, the data cables are installed in a carrier constructed of electrical metallic tubing (EMT), ferrous conduit or pipe, or rigid sheet steel ducting. All connections in a hardened system are permanently sealed around all surfaces with welds, epoxy, or other such sealants. If the hardened carrier is buried underground, to secure cables running between buildings, for example, the carrier containing the cables is encased in concrete.

Secure Areas

Secure areas are sites or spaces where you handle sensitive information or shelter/store valuable IT equipment and personnel to achieve business objectives. In the context of physical security, the term site means buildings, rooms, or offices that host all the services and facilities (electricity, heating, air conditioning). This falls under ISO 27001, which involves physical security and keeping secure areas secure. The following components of secure areas help enhance the capabilities of these areas:



- **Air gap:** As previously mentioned, air-gapped computer systems have no physical connection to other computers, networks, or unsecured systems. Why would you need an air-gapped computer? Say you want to work on sensitive business documents or financials without the risk of malware, ransomware, or keyloggers (any attacks done over Internet networks). In this case, you can simply set up an offline computer that is not connected to the Internet.
- **Vault:** A bank vault is a secure space to store money, valuables, records, and documents. In a secure space, a vault is intended to protect contents from theft, unauthorized use, fire, natural disasters, and other threats, much like a safe. Vaults are typically much larger than safes and can even house entire data centers.
- **Safe:** A safe is a device or type of housing that is capable of storing valuable information, such as files, tapes, money, and the like. It is intended for short-term storage of company assets. A fireproof safe is recommended for business because it is capable of withstanding small fires.

- ***Hot aisle*:** A hot aisle/cold aisle is a layout design for server racks and other computing equipment in a data center. The goal of a hot aisle/cold aisle configuration is to conserve energy and lower cooling costs by managing air flow. In its simplest form, hot aisle/cold aisle data center design involves lining up server racks in alternating rows with cold air intakes facing one way and hot air exhausts facing the other. The rows that the heated exhausts pour into are called hot aisles. Typically, hot aisles face air conditioner return ducts.
- ***Cold aisle*:** The rows composed of rack fronts are called cold aisles. Typically, cold aisles face air conditioner output ducts. A containment system can be used to isolate hot aisles and cold aisles from each other and prevent hot and cold air from mixing. Containment systems started out as physical barriers that simply separated the hot and cold aisles with vinyl plastic sheeting or Plexiglas covers; now they are almost like complete rooms enclosing the hot aisle and cold aisle.

Secure Data Destruction

An organization's security starts at the purchasing of equipment, systems, and services and ends with the removal and destruction of those items. You may have heard of someone purchasing a server off the Internet only to find the original owner didn't erase the data. Companies must create and put into practice a data destruction policy. It is not enough to erase data from hard drives of servers or workstations; recovery programs can still get at that data. A "shred-all" policy is one of the most critical steps you can take toward total information security; it means a company-wide commitment to shredding all data and documents on a regular basis. Standardizing data and document destruction procedures enables your organization to align its governance policy and regulations with its information security goals and needs:

Key Topic

- **Burning:** One of the best methods in documentation destruction is fire. Without taking very extreme measures, recovering data from ashes is nearly impossible.
- **Shredding:** One of the most cost-effective methods for bulk paper destruction is shredding. Shredders are able to malform paper into multiple strips with jagged edges. Similarly, hard drives and other magnetic media can be shredded via large industrial shredders.
- ***Pulping*:** Paper also can be shredded and then reduced to pulp. After the paper has been broken down into pulp, the fibers are separated, washed, and screened to remove fiber bundles. Then the water is pressed out. After the pulp is dried, it can be made into recycled paper. Because all the ink letters from the paper have been removed, documents cannot be recovered.

- ***Pulverizing*:** Pulverizing grinds and/or reshreds paper to make it impossible to recover. Documents and storage media are fed into a pulverizer that uses hydraulic or pneumatic action to reduce the materials to loose fibers and shards. One disadvantage to this service is that it is very costly, and only a few commercial disposers use this method; however, it does guarantee complete data destruction.
- ***Degaussing*:** Degaussing is the process of reducing or eliminating a magnetic field (or data) stored on tape and disk media such as computer and laptop hard drives, disks, reels, cassettes, and cartridge tapes. By using the right degausser, you can guarantee that your information will no longer be retrievable. A degausser creates a magnetic field of an intensity that erases and corrupts the layout/form of the data on magnetic media.
- **Third-party solutions:** There are a number of third-party secure document and electronic media destruction and disposal services. They typically have large facilities and commercial machinery to manage large volumes of documents and media. These organizations securely transport your documents to their facilities, where they can shred, pulp, or pulverize. Depending on your contract, you may receive large trash receptacles where people can drop their documents. These containers are locked and are picked up and exchanged on a regular basis.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 15-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 15-2 Key Topics for Chapter 15

Key Topic Element	Description	Page Number
Section	Bollards/Barricades	370
Section	Access Control Vestibules	372
Section	Badges	373
Section	Alarms	374

Key Topic Element	Description	Page Number
Section	Signage	374
Section	Cameras	375
Section	Closed-Circuit Television (CCTV)	376
Section	Industrial Camouflage	377
List	Personnel	377
List	Locks	378
Section	USB Data Blockers	379
Section	Lighting	380
Section	Fencing	380
Section	Fire Suppression	381
List	Sensors	382
Section	Drones	382
Section	Visitor Logs	383
Section	Faraday Cages	383
Section	Air Gap	384
Section	Screened Subnet (Previously Known as Demilitarized Zone [DMZ])	384
Section	Protected Cable Distribution	385
List	Secure areas	385
List	Secure data destruction	386

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

bollard, barricade, access control vestibules, motion recognition, object detection, closed-circuit television (CCTV), industrial camouflage, guards, robot sentries, reception, two-person integrity/control, biometrics:, electronic locks, physical locks, cable locks, USB data blocker, fire suppression, Faraday cage, air gap, screened subnet, protected cable distribution system (PDS), hot aisle, cold aisle, pulping, pulverizing, degaussing

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What are badges used for in physical security controls?
2. What purpose does signage serve in controlling security in a building or factory?
3. What does industrial camouflage accomplish in today's business and industrial environment?
4. How does two-person integrity/control ensure systems and corporate data integrity are accomplished?
5. What sensors provide access to an area?



This chapter covers the following topics related to Objective 2.8 (Summarize the basics of cryptographic concepts) of the CompTIA Security+ SY0-601 certification exam:

- Digital signatures
- Key length
- Key stretching
- Salting
- Hashing
- Key exchange
- Elliptic-curve cryptography
- Perfect Forward Secrecy
- Quantum
  - Communications
  - Computing
- Post-quantum
- Ephemeral
- Modes of operation
  - Authenticated
  - Unauthenticated
  - Counter
- Blockchain
  - Public ledgers
- Cipher suites
  - Stream
  - Block
- Symmetric vs. asymmetric
- Lightweight cryptography
- Steganography
  - Audio
  - Video
  - Image
- Homomorphic encryption
- Common use cases
  - Low power devices
  - Low latency
  - High resiliency
  - Supporting confidentiality
  - Supporting integrity
  - Supporting obfuscation
  - Supporting authentication
  - Supporting non-repudiation
- Limitations
  - Speed
  - Size
  - Weak keys
  - Time
  - Longevity
  - Predictability
  - Reuse
  - Entropy
  - Computation overheads
  - Resource vs. security constraints

Summarizing the Basics of Cryptographic Concepts

Cryptography is the art and science of keeping data secure. Cryptographic techniques are used to encrypt (scramble) messages before the messages are stored or transmitted. When data is encrypted, the message and an encryption key are passed to the encryption algorithm. To decrypt the data, the ciphertext and a decryption key are passed to the decryption algorithm.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 16-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 16-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Digital Signatures	1
Hashing	2
Quantum	3–4
Blockchain	5
Cipher Suites	6
Steganography	7
Common Use Cases	8
Limitations	9
Key Stretching	10
Key Derivation	11

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which is one of the three digital signature scheme algorithms?
   a. A key generation algorithm that selects a private key
   b. A key generation algorithm that selects a data path
   c. A file encryption algorithm that encrypts files
   d. A triple-DES communication encryption algorithm

2. Which of the following properties are not ideal cryptographic hash function properties?
   a. Provides a quick response to any mathematical subfunction
   b. Is unable to compute the hash value for any given message
   c. Is quick to compute the cryptographic key's salt value
   d. Is infeasible to find two different messages with the same hash value

3. The complexity of quantum properties lies in the principles of quantum mechanics that are behind quantum cryptography. Which is one of the primary principles?
   a. Changing the zero bit in the *x* state moves the bit 35 degrees left.
   b. Measuring the quantum property adds two qubits.
   c. Quantum mechanics cannot fix a laser while in operation.
   d. You can't measure a quantum property without changing or disturbing it.

4. When a sender transmits photons through a filter that randomly gives one of four possible polarizations and bit designations, what is it set as?
   a. 35-degree left (zero bit)
   b. Horizontal (zero bit)
   c. Vertical (45 bit)
   d. Top 45-degree left (seven bit)

5. A blockchain is a specific type of database that differs from a typical database in the way it stores information. Which is the most accurate description of storage in a blockchain?
   a. Data is stored in blocks that are then chained together.
   b. Data is stored in segments that are linked.
   c. Data is linked to a SQL database stored in blocks.
   d. Data is chained to specific blocks in a linked mathematical code.

6. Today's cipher suite is a set of algorithms that help secure a network connection that uses Transport Layer Security (TLS). The set of algorithms that cipher suites usually contain includes which of the following?
   a. A bulk repository hashing key
   b. A bulk encryption algorithm
   c. A hash exchange and management algorithm
   d. Key exchange repository for storage

7. Steganography is the practice of hiding a secret message inside or even on top of something that is not secret. Which of the following is not a tool that can be used to perform steganography on data?
   a. Steghide
   b. Foremost
   c. Stegais
   d. Future

8. IoT low-power devices have limited power. What effect does that have on a cryptographic algorithm that does not compromise security?
   a. Smaller symmetric key sizes
   b. The use of dual asymmetric keys
   c. The use of 32-bit 3DES encryption
   d. The use of shorter RC4 encryption keys

9. Timing attacks that use side channels to analyze the time taken to execute a cryptographic algorithm are based on which method?
   a. Smaller keys for quicker analysis
   b. Larger keys for faster mathematical calculations
   c. Time
   d. Length

10. Which of the following runs a password through an algorithm to produce an enhanced key, usually at least 128 bits long?
    a. IV
    b. Steganography
    c. Entropy
    d. Key stretching

11. Which of the following are key derivation functions (KDFs) that are primarily used for key stretching? (Select two.)
    a. Symmetric key algorithms
    b. Asymmetric key algorithms
    c. Bcrypt
    d. PBKDF2

Foundation Topics

Digital Signatures

Key Topic

Digital signatures are a mathematical scheme for verifying the authenticity of digital messages or documents. A valid digital signature, where the prerequisites are satisfied, gives a recipient very strong reason to believe that the message was created by a known sender (authentication) and that the message was not altered in transit (integrity).

Digital signatures are a standard element of most cryptographic protocol suites and are commonly used for software distribution, financial transactions, contract management software, and in other cases where it is important to detect forgery or tampering. Digital signatures are often used to implement electronic signatures, which includes any electronic data that carries the intent of a signature.

Digital signatures employ asymmetric cryptography. In many instances, they provide a layer of validation and security to messages sent through a nonsecure channel. Properly implemented, a digital signature gives the receiver reason to believe the message was sent by the claimed sender. Digital signatures properly implemented are more difficult to forge. Digital signature schemes, in the sense used here, are cryptographically based and must be implemented properly to be effective. Digital signatures can also provide nonrepudiation, meaning that the signer cannot successfully claim he or she did not sign a message, while also claiming his or her private key remains secret. Some nonrepudiation schemes offer a timestamp for the digital signature so that even if the private key is exposed, the signature is valid.

A digital signature scheme typically consists of three algorithms:

- A key generation algorithm that selects a private key uniformly at random from a set of possible private keys. The algorithm outputs the private key and a corresponding public key.
- A signing algorithm that, given a message and a private key, produces a signature.
- A signature verifying algorithm that, given the message, public key, and signature, either accepts or rejects the message's claim to authenticity.

Two main properties are required for a digital signature:

- The authenticity of a signature generated from a fixed message and fixed private key can be verified by using the corresponding public key.
- Generating a valid signature for a party without knowing that party's private key should be computationally infeasible.

A digital signature is an authentication mechanism that enables the creator of the message to attach a code that acts as a signature. The Digital Signature Algorithm (DSA), developed by the National Institute of Standards and Technology (NIST), is one of many examples of a signing algorithm.

NOTE As organizations move away from paper documents with ink signatures, digital signatures can provide added assurances of the evidence to provenance, identity, and status of an electronic document and also acknowledge informed consent and approval by a signatory.



Key Length

Key length is equal to the number of bits in an encryption algorithm's key. A short key length means poor security. However, a long key length does not necessarily mean good security. The key length determines the maximum number of combinations required to break an encryption algorithm.

The strength of encryption is related to the difficulty of discovering the key, which in turn depends on both the cipher used and the length of the key.

Encryption strength is often described in terms of the size of the keys used to perform the encryption: longer keys may provide stronger encryption. Key length is measured in bits. The 128-bit keys for use with the RC4 symmetric-key cipher supported by SSL provide significantly better cryptographic protection than 40-bit keys for use with the same cipher. Obviously, 128-bit RC4 encryption is 3×1026 times stronger than 40-bit RC4 encryption.

Different ciphers require different key lengths to achieve the same level of encryption strength. The RSA cipher used for public-key encryption, for example, can use only a subset of all possible values for a key of a given length, due to the nature of the mathematical problem on which it is based. Other ciphers, such as those used for symmetric key encryption, can use all possible values for a key of a given length, rather than a subset of those values.

Thus, a 128-bit key for use with a symmetric-key encryption cipher would provide stronger encryption than a 128-bit key for use with the RSA public-key encryption cipher. This difference explains why the RSA public-key encryption cipher must use a 512-bit key (or longer) to be considered cryptographically strong, whereas symmetric key ciphers can achieve approximately the same level of strength with a 64-bit key.

NOTE Many cryptographical implementations are broken before the actual encryption itself is broken; see RC4 and WEP.

Key Topic

Key Stretching

Key stretching techniques are used to make a possibly weak key, typically a password or passphrase, more secure against a brute-force attack by increasing the resources (time and possibly space) needed to test each possible key. Key stretching is the practice of converting a password to a longer and more random key for cryptographic purposes. The process of converting a password into a key is accomplished by a type of algorithm known as a key derivation function that may include salt and pepper with the password to make the key more difficult to guess. Table 16-2 provides an overview of key stretching, highlighting its strengths and weaknesses.

Table 16-2 Key Stretching

Type	Cryptography Information Security
Definition	The practice of converting a password to a longer and more random key for cryptographic purposes such as encryption.
Notes	Key stretching can also be applied to a master key as opposed to a password. If a password is extremely long, key stretching can result in a shorter and less secure key.
Related Concepts	Cryptography Key Derivation Function Strong Password Cryptographic Keys Information Security
Salt	Generation of random data that is used as an additional input to a one-way function that hashes data.
Pepper	Like a salt in that it is a randomized value added to a password hash, and it should be kept secret. Also referred to as a secret salt.

Key Topic

Salting

In cryptography, ***salting*** refers to the generation of random data that is used as an additional input to a one-way function that hashes data, a password, or a passphrase. Salts are used to safeguard passwords in storage. A new salt is randomly generated

for each password. In a typical setting, the salt and the password (or its version after key stretching) are concatenated and processed with a cryptographic hash function, and the output hash value (but not the original password) is stored with the salt in a database.

Salts defend against a precomputed hash attack. Because salts are different in each case, they also protect commonly used passwords, or those users who use the same password on several sites, by making all salted hash instances for the same password different from each other. Figure 16-1 depicts the password salting process.

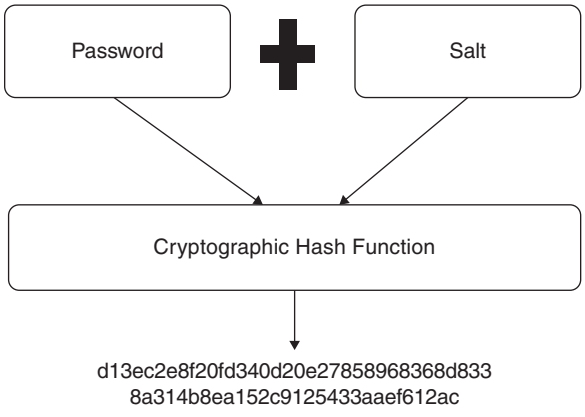


FIGURE 16-1 Password Salting

Key Topic

Hashing

A cryptographic ***hash*** function is a mathematical algorithm that maps data of arbitrary size to a bit array of a fixed size. It is a one-way function—that is, a function that is practically infeasible to invert. Ideally, the only way to find a message that produces a given hash is to attempt a brute-force search of possible inputs to see if they produce a match or use a rainbow table of matched hashes.

The ideal cryptographic hash function has the following main properties:

- It is deterministic, meaning that the same message always results in the same hash.
- It is quick to compute the hash value for any given message.
- It is infeasible to generate a message that yields a given hash value or to reverse the process that generated the given hash value.
- It is infeasible to find two different messages with the same hash value.

- A small change to a message should change the hash value so extensively that the new hash value appears uncorrelated with the old hash value avalanche effect.

Cryptographic hash functions have many information security applications, notably in digital signatures, message authentication codes (MACs), and other forms of authentication. They can also be used as ordinary hash functions, to index data in hash tables, for fingerprinting, to detect duplicate data or uniquely identify files, and as checksums to detect accidental data corruption.

Key Topic

Key Exchange

Key exchange, also known as key establishment, is a method in cryptography by which cryptographic keys are exchanged between two parties, allowing use of a cryptographic algorithm. If the sender and receiver wish to exchange encrypted messages, each must be equipped to encrypt messages to be sent and to decrypt messages received. The nature of the equipping they require depends on the encryption technique they might use. If they use a code, both require a copy of the same source code. If they use a cipher, they need appropriate keys. If the cipher is a symmetric key cipher, both need a copy of the same key. If it is an asymmetric key cipher with the public/private key property, both need the other's public key. Figure 16-2 illustrates the key exchange process.

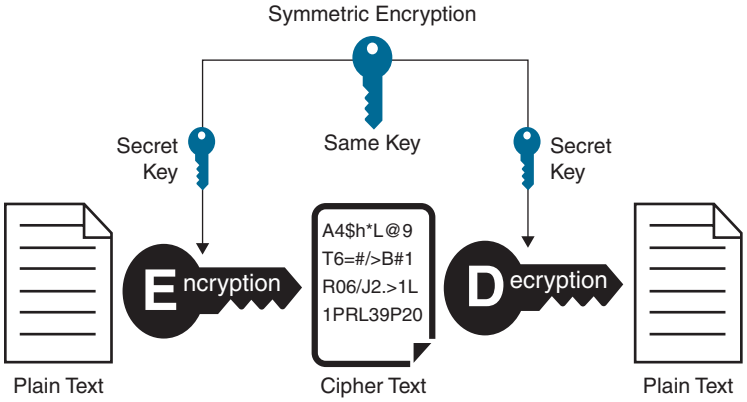


FIGURE 16-2 Symmetric Key Exchange Process

Key Topic

Elliptic-Curve Cryptography

Elliptic-curve cryptography (ECC) is an approach to public-key cryptography based on the algebraic structure of elliptic curves over finite fields. ECC allows smaller keys compared to non-EC cryptography based on plain Galois fields to provide

equivalent security. An algebraic structure consists of a nonempty set "A" (called the underlying set, carrier set, or domain); a collection of operations on A of finite arity, typically binary operations; and a finite set of identities, known as axioms, that these operations must satisfy.

Elliptic curves are applicable for key agreement and digital signatures, and indirectly can be used for encryption by combining the key agreement with a symmetric encryption scheme. They are also used in several integer factorization algorithms based on elliptic curves that have applications in cryptography, such as Lenstra elliptic-curve factorization.

An elliptic curve is a plane curve over a finite field, rather than the real numbers that consist of the points satisfying the following equation:

$$y^2 = x^3 + ax + b$$

To use ECC, all parties must agree on all the elements defining the elliptic curve, the domain parameters of the scheme. The domain parameters are not typically generated by each participant. Calculating the number of points on a curve is time consuming and complicated to implement. As a result, several standards bodies published domain parameters of elliptic curves for several common field sizes. Such domain parameters are commonly known as standard curves or named curves; a named curve can be referenced either by name or by the unique object identifier.

Perfect Forward Secrecy

Key Topic

In cryptography, forward secrecy, also known as ***perfect forward secrecy***, is a feature of specific key agreement protocols that gives assurances that session keys will not be compromised even if long-term secrets used in the session key exchange are compromised.

Perfect forward secrecy means that a piece of an encryption system automatically and frequently changes the keys it uses to encrypt and decrypt information, such that if the latest key is compromised, it exposes only a small portion of the user's sensitive data.

Encryption tools with perfect forward secrecy switch their keys as frequently as every message in text-based conversation, every phone call in the case of encrypted calling apps, or every time a user loads or reloads an encrypted web page in his or her browser.

While schemes for perfect forward secrecy date back to the early 1990s, the feature was first practically implemented in Off-The-Record Messaging, a protocol for encrypted instant messaging invented in 2004 that encrypted messages with a new

key every time a sender alternated in an instant messaging conversation. In that system, multiple messages sent back to back by the same sender still used the same key.



Quantum

Quantum cryptography is the science of exploiting quantum mechanical properties to perform cryptographic tasks. The best-known example of quantum cryptography is quantum key distribution, which offers an information-theoretically secure solution to the key exchange problem.

Cryptography is the process of encrypting data or converting plain text into scrambled text so that only someone who has the right "key" can decode and read it. Quantum cryptography, by extension, simply uses the principles of quantum mechanics to encrypt data and transmit it in a way that cannot be hacked.

While the definition sounds simple, the complexity lies in the principles of quantum mechanics behind quantum cryptography:

- The particles that make up the universe are inherently uncertain and can simultaneously exist in more than one place or more than one state of being.
- Photons are generated randomly in one of two quantum states.
- You can't measure a quantum property without changing or disturbing it.
- You can clone some quantum properties of a particle, but not the whole particle.

Communications

Quantum cryptography, or quantum key distribution (QKD), uses a series of photons (light particles) to transmit data from one location to another over a fiber-optic cable. By comparing measurements of the properties of a fraction of these photons, the two endpoints can determine what the key is and if it is safe to use. Breaking the process down further helps to explain it better.

The sender transmits photons through a filter (or polarizer), which randomly gives them one of four possible polarizations and bit designations: vertical (one bit), horizontal (zero bit), 45-degree right (one bit), or 45-degree left (zero bit).

The photons travel to a receiver, which uses two beam splitters (horizontal/vertical and diagonal) to "read" the polarization of each photon. The receiver does not know which beam splitter to use for each photon and has to guess which one to use.

Once the stream of photons has been sent, the receiver tells the sender which beam splitter was used for each of the photons in the sequence they were sent, and the sender compares that information with the sequence of polarizers used to send

the key. The photons that were read using the wrong beam splitter are discarded, and the resulting sequence of bits becomes the key.

If the photon is read or copied in any way by an eavesdropper, the photon's state changes. The change is detected by the endpoints. In other words, this means you cannot read the photon and forward it on or make a copy of it without being detected.

Computing

Quantum computing is the use of quantum phenomena such as superposition and entanglement to perform computation. Computers that perform quantum computations are known as quantum computers.

Post-Quantum

Key Topic

Post-quantum cryptography refers to cryptographic algorithms (usually public-key algorithms) that are thought to be secure against an attack by a quantum computer. These complex mathematical equations take traditional computers months or even years to break. If you are planning to store "secret" or private data for long periods of time, how can you protect it—for the long term? This is where post-quantum approaches can be implemented using today's computers, but that would be impervious to attacks from tomorrow's quantum units. When you increase the size of digital keys, the number of permutations that need to be searched using brute computing power rises exponentially. Just doubling the size of a key from 128 bits to 256 bits effectively squares the number of possible permutations that a quantum machine using Grover's algorithm would have to search through. Grover's algorithm is a quantum algorithm that finds with high probability the unique input to a black box function that produces a particular output value.

By contrast, quantum computers running Shor's algorithm can break math-based systems in moments. Shor's algorithm is a polynomial-time quantum computer algorithm for integer factorization. Informally, it solves the following problem: Given an integer N, find its prime factors.

NIST is reviewing and working on 26 post-quantum proposals for a government standard because there is currently no set standard; however, one is expected to be released in draft sometime in 2022.

Quantum cryptography uses the principles of quantum mechanics to secure messages and, unlike mathematical encryption, is truly unhackable.

Key Topic

Ephemeral

A cryptographic key is called ephemeral if it is generated for each execution of a key-establishment process and meets other requirements of the key type, unique to each message or session.

Ephemeral key agreement keys are the private key plus its corresponding public keys of asymmetric key pairs that are used for a single key establishment transaction to establish one or more keys and optionally other keying material such as initialization vectors (IVs).

Ephemeral keys are designed to be used for a single transaction or session. *Ephemeral* describes a temporary or short duration; in this case it's a key generated for a single use. Figure 16-3 illustrates the ephemeral key process.

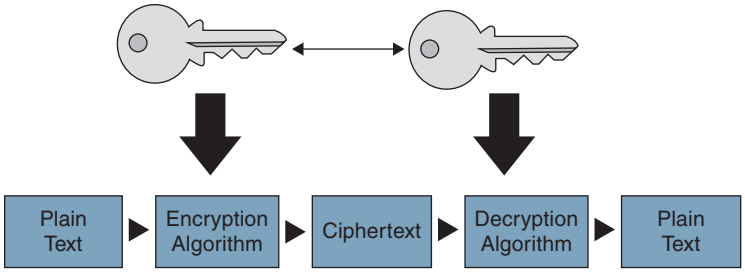


FIGURE 16-3 Ephemeral Key Process

Key Topic

Modes of Operation

A mode of operation describes how to repeatedly apply a cipher's single-block operation to securely transform amounts of data larger than a block. Most modes require a unique binary sequence, often called an initialization vector, for each encryption operation.

The IV has to be nonrepeating and, for some modes, random as well. The initialization vector is used to ensure distinct ciphertexts are produced even when the same plaintext is encrypted multiple times independently with the same key. Block ciphers might be capable of operating on more than one block size, but during transformation, the block size is always fixed. Block cipher modes operate on whole blocks and require that the last part of the data be padded to a full block if it is smaller than the current block size. There are, however, modes that do not require padding because they effectively use a block cipher as a stream cipher.

Later cipher development regarded integrity protection as an entirely separate cryptographic goal. Some modern modes of operation combine confidentiality and authenticity in an efficient way and are known as authenticated encryption modes.

Authenticated modes include authenticated encryption (AE) and authenticated encryption with additional or associated data (AEAD). With AE, the requirement is to both protect the privacy of the message and to ensure authenticity. A method for achieving both of these goals at the same time is called authenticated encryption. With AEAD, the header needs to be authenticated but should not be encrypted. The task of encrypting the message and authenticating both the header and the message is called authenticated encryption with associated data. There are other methods such as deterministic authenticated encryption with associated data (DAEAD). AE/AD constructions use a nonce, which is distinct for every message. Deterministic authenticated encryption does away with the nonce. Only the message is processed using a secret (random) key to produce the ciphertext. An extension of this functionality allows the authentication of associated data and the message.

In ***unauthenticated mode***, communication is encrypted, but users do not have to authenticate or verify their credentials. Deployment automation automatically uses this mode for server/agent communication. SSL unauthenticated mode can also be used for HTTP communication.

Counter mode uses an arbitrary number, the counter, that changes with each block of text encrypted. The counter is encrypted with the cipher, and the result is XOR'd (exclusive OR'd) into ciphertext. Because the counter changes for each block, the problem of repeating ciphertext that results from the Electronic Code Book method is avoided.

Electronic Code Book Mode

Electronic Code Book (ECB) is a mode of operation for a block cipher, with the characteristic that each possible block of plaintext has a defined corresponding ciphertext value. In fact, ECB can support a separate encryption key for each block type. ECB encodes each block independently. A problem arises when there is identical input; in this case, you end up with identical output, so it is predictable. Figure 16-4 shows this function in action.

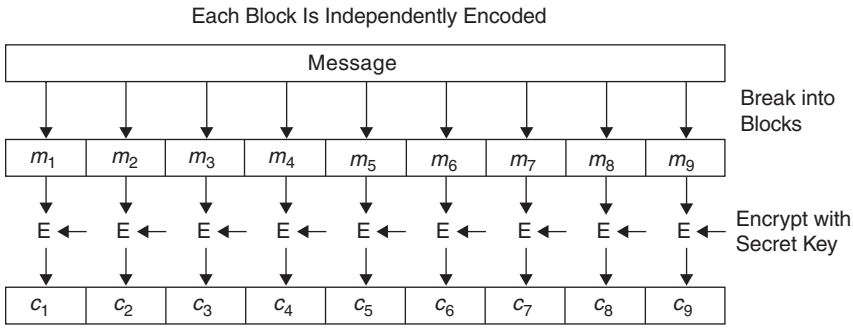


FIGURE 16-4 ECB Electronic Code Book

Cipher Block Chaining Mode

Cipher Block Chaining (CBC) is a mode of operation for a block cipher, one in which a sequence of bits is encrypted as a single unit or block with a cipher key applied to the entire block. The initialization vector is of a certain length. One of its key characteristics is that it uses a chaining mechanism that causes the decryption of a ciphertext block to depend on all the preceding ciphertext blocks. One limitation of CBC mode is that encryption errors propagate: an encryption error in one block cascades through subsequent blocks due to the chaining, destroying their integrity.

Rearrangement of the order of the ciphertext blocks causes decryption to become corrupted. Basically, in cipher block chaining, each plaintext block is XOR'd with the immediately previous ciphertext block and then encrypted. Identical ciphertext blocks can result only if the same plaintext block is encrypted using both the same key and initialization vector, and if the ciphertext block order is not changed. This mode has an advantage over Electronic Code Book mode in that the XOR'ing process hides plaintext patterns. Figure 16-5 shows this function.

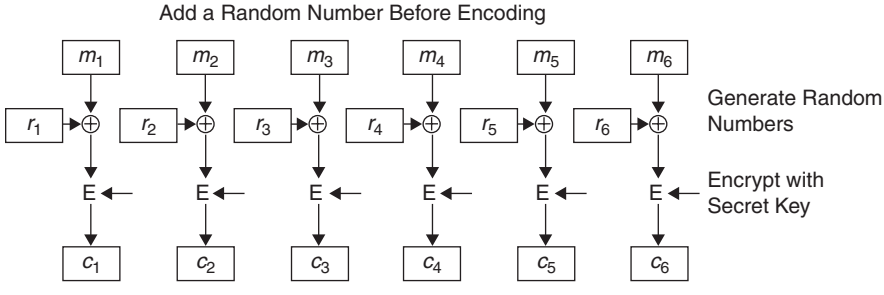


FIGURE 16-5 Cipher Block Chaining (CBC)

Cipher Feedback Mode

Cipher Feedback (CFB) mode is similar to CBC mode, but the primary difference is that CFB is a stream mode cipher. Like CBC, CFB uses an initialization vector and destroys patterns, so errors will be propagated.

CFB mode allows a block cipher with a block size of *n* bits to be used as a stream cipher with a data encryption unit of *m* bits. This is represented as for any $m \leq n$. For example, to encrypt *m* bits of data, the block cipher is used to encrypt the contents of the register, the left most *m* bits of the result are XOR'd with the *m* bits of data, and the result is *m* bits of ciphertext. In addition, the register is shifted left by *m* bits, and those *m* ciphertext bits are inserted in the rightmost *m* register bits to be used in processing the next *m* bits of plaintext. To run the CFB operation mode as a stream cipher, FIPS 81 defines CFB variations where plaintext blocks can have any size less than 64 bits. To describe CFB variations, you need the following additional notations:

Decryption reverses the process. The register initially contains the initialization vector. To decrypt *m* bits of ciphertext, the block cipher is used to encrypt the contents of the register, and the resulting leftmost *m* bits are XOR'd with the *m* ciphertext bits to recover *m* plaintext bits. The *m* ciphertext bits are then shifted left into the register.

Unlike in CBC, a single transmission error in the ciphertext causes an error in $n/m + 1$ blocks as the affected ciphertext block is shifted through the register and then the system recovers. With 8-bit CFB, if a byte is lost, one byte of plaintext is lost and the next 8 bytes are garbled. After that, the plaintext decrypts properly. If a byte is added to the ciphertext, a byte of garbage is added, and the following 8 bytes are garbled but the rest is okay.

NOTE The encryption function of the block cipher is used in encryption and decryption of CFB mode, and the decryption function of the block cipher is not used at all.

Figure 16-6 shows the *k*-bit stream block function and how *k*-bit functions operate in CFB mode.

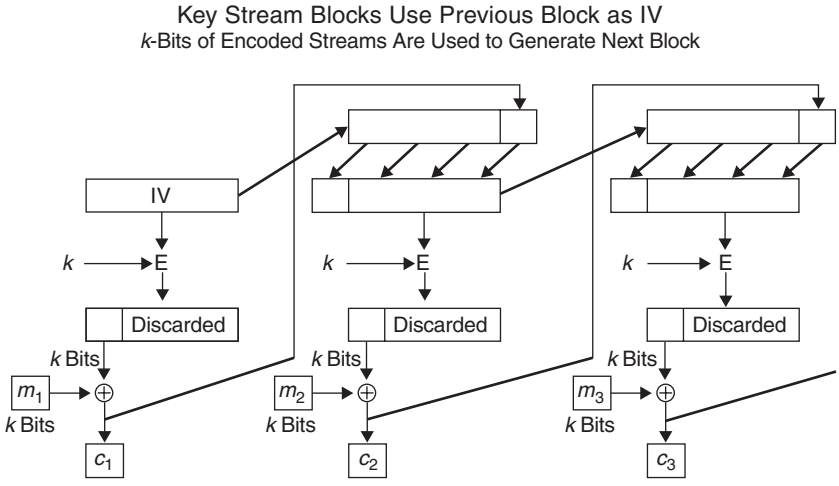


FIGURE 16-6 *k*-bit Cipher Feedback (CFB) Mode

Output Feedback Mode

Output Feedback (OFB) mode is a mode of operation for a block cipher. It has some similarities to the Ciphertext Feedback mode in that it permits encryption of differing block sizes, but a key difference is that the output of the encryption block function is the feedback (instead of the ciphertext). The XOR (exclusive OR) value of each plaintext block is created independently of both the plaintext and ciphertext. This mode is used when there can be no tolerance for error propagation because there are no chaining dependencies. Like the ciphertext feedback mode, it uses an initialization vector. Changing the IV in the same plaintext block results in different ciphertext.

In terms of error correction, output feedback can tolerate ciphertext bit errors but is incapable of self-synchronization after losing ciphertext bits because it disturbs the synchronization of the aligning keystream. A problem with output feedback is that the plaintext can be easily altered, but by using a digital signature scheme, you can overcome this problem. Figure 16-7 illustrates OFB mode.

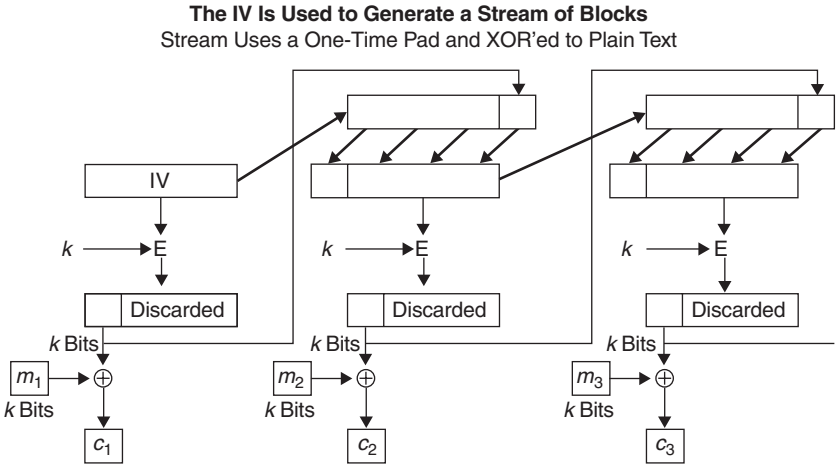


FIGURE 16-7 *k*-Bit OFB Mode

Key Topic

Counter Mode

Counter (CTR) mode is a simple counter-based block cipher implementation in cryptography. Every time a counter-initiated value is encrypted and given as input to XOR with plaintext or original text, it results in a ciphertext block. The CTR mode is independent of feedback use and thus can be implemented in parallel in this mode. CTR generates the next keystream block by encrypting successive values named counters.

This counter can be any purpose or function which generates a sequence that is guaranteed not to call for a long time, although an actual increment-by-one counter is the simplest and most popular. When utilization of the easiest deterministic input function, the implementation reveals a cryptosystem with known systematic input, and therefore represents a risk.

CTR mode, like OFB, turns a block cipher into a stream cipher. It generates the next keystream block by encrypting successive values of a counter. CTR mode has similar characteristics to OFB but also allows a random-access property during decryption.

CTR can be considered as a counter-based version of CFB mode without the feedback in this cryptography. In this mode, both the sender and receiver need to get access to a reliable counter, which finds out a new shared value each time a ciphertext block is exchanged or changes the value, which is shared. This shared counter is not necessarily a secret value, but the challenge is that both sides must keep the counter synchronized at a time when both are active. Figure 16-8 illustrates CTR mode.

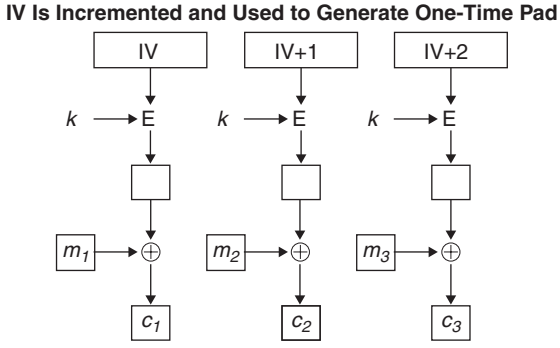


FIGURE 16-8 Counter (CTR) Mode

The main disadvantage of CTR mode is that it requires a synchronous counter at the sender and receiver. Decreases in synchronization lead to the incorrect recovery of plaintext or original text.

Key Topic

Blockchain

Blockchain is a specific type of database. It differs from a typical database in the way it stores information: blockchains store data in blocks that are then chained together. As new data comes in, it is entered into a fresh block. When a block is filled with data, it is chained onto the previous block, which makes the data chained together in chronological order. Different types of information can be stored on a blockchain, but the most common use so far has been as a ledger for transactions.

Blockchain is a system of recording information in a way that makes it difficult or impossible to change, hack, or cheat the system.

A blockchain is essentially a digital ledger of transactions that is duplicated and distributed across the entire network of computer systems on the blockchain. Each block in the chain contains a number of transactions, and every time a new transaction occurs on the blockchain, a record of that transaction is added to every participant's ledger. The decentralized database managed by multiple participants is known as Distributed Ledger Technology (DLT).

The ***public ledger*** organizes into a long chain of blocks of information. When a buyer and seller engage in a transaction, the blockchain verifies the authenticity of their accounts. This is done by using the public ledger and by checking if the funds are available to proceed with the transactions. However, if the funds are not available in the buyer's account or are promised to another party, the sale is prevented, effectively making double buying impossible. Instead of the public ledger being maintained and monitored by a central authority, like banks or governments, the public

ledger is stored on the personal computers or electronic devices of the individuals and businesses who use the blockchain. To use the blockchain, you must donate to the collective. In this way, the need for a central authority is eliminated because every member of the blockchain has access to the ledger. Further, the ledger itself cannot be falsified. If one individual tampers with his or her ledger, the blockchain verifies that ledger against the other ledgers owned by the rest of the community and rejects it quickly.



Cipher Suites

A ***cipher suite*** is a set of algorithms that help secure a network connection that uses Transport Layer Security (TLS). Cipher suites usually contain a set of algorithms that include a key exchange algorithm, bulk encryption algorithm, and message authentication code (MAC) algorithm.

The key exchange algorithm is used to exchange a key between two devices. This key is used to encrypt and decrypt the messages being sent between two machines. The bulk encryption algorithm is used to encrypt the data being sent. The MAC algorithm provides data integrity checks to ensure that the data sent does not change in transit. In addition, cipher suites can include signatures and an authentication algorithm to help authenticate the server and/or client. Hundreds of different cipher suites contain different combinations of these algorithms. Some cipher suites offer better security than others.

The structure and use of the cipher suite concept are defined in the TLS standard document. TLS 1.2 was the most prevalent version of TLS. The newer version of TLS (TLS 1.3) includes additional security requirements to cipher suites. Cipher suites defined for TLS 1.2 cannot be used in TLS 1.3, and vice versa, unless otherwise stated in their definition. TLS 1.2 was defined in RFC 5246 and was in use after 2008 by the majority of all web browsers. On March 21, 2018, TLS 1.3 was finalized. As of August 2018, the final version of TLS 1.3 was published (RFC 8446). In short, the major benefits of TLS 1.3 versus that of TLS 1.2 are improved security and faster speed.

A ***stream cipher*** is a symmetric key cipher in which plaintext digits are combined with a pseudorandom cipher digit stream known as a keystream. In a stream cipher, each plaintext digit is encrypted one at a time with the corresponding digit of the keystream, to give a digit of the ciphertext stream. Because encryption of each digit is dependent on the current state of the cipher, it is also known as a state cipher. A digit is typically a bit, and the combining operation is an exclusive OR (XOR).

The pseudorandom keystream is typically generated serially from a random seed value using digital shift registers. The seed value serves as the cryptographic key for decrypting the ciphertext stream. Stream ciphers represent a different approach to symmetric encryption from block ciphers.

Stream ciphers typically execute at a higher speed than block ciphers and have lower hardware complexity. However, stream ciphers can be susceptible to serious security problems if used incorrectly; in particular, the same starting state seed must never be used twice.

A ***block cipher*** is an encryption method that applies a deterministic algorithm along with a symmetric key to encrypt a block of text instead of encrypting one bit at a time as in stream ciphers. For example, a common block cipher, AES, encrypts 128-bit blocks with a key of predetermined length: 128, 192, or 256 bits. Block ciphers operate on large blocks of digits with a fixed, unvarying transformation. This distinction is not always clear-cut. In some modes of operation, a block cipher primitive is used in such a way that it acts effectively as a stream cipher.

Key Topic

Symmetric vs. Asymmetric Encryption

Asymmetric encryption is more secure than symmetric encryption. ***Symmetric encryption*** uses a single key that needs to be shared among the people who need to receive the message, whereas ***asymmetric encryption*** uses a pair of public keys and a private key to encrypt and decrypt messages when communicating. Asymmetric algorithms are often referred to as *public key algorithms* because they use the public key as the focal point for the algorithm.

As shown in Figure 16-9, symmetric is the simplest kind of encryption; it involves only one secret key to cipher and decipher information. Symmetric encryption is an old and the best-known technique. It uses a secret key that can either be a number, word, or string of random letters. It is a blended with the plaintext of a message to change the content in a particular way. The sender and recipient should know the secret key that is used to encrypt and decrypt all the messages.

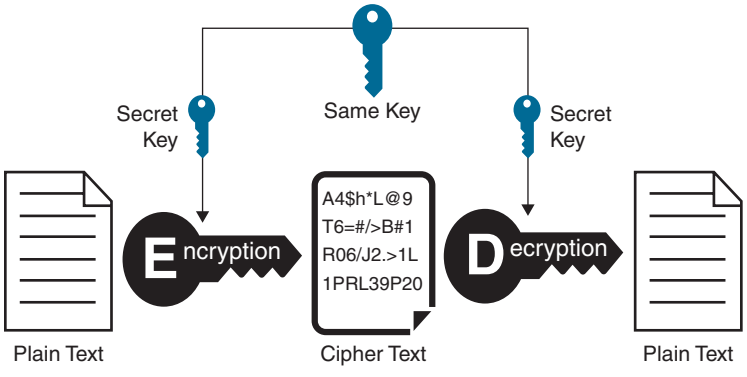


FIGURE 16-9 Symmetric Encryption

Examples of symmetric encryption include Blowfish, Twofish, AES, RC4, and DES. The most widely used symmetric algorithms are AES-128, AES-192, and AES-256. The main disadvantage of symmetric key encryption is that all parties involved have to exchange the key used to encrypt the data before they can decrypt it.

Blowfish is a symmetric-key block cipher. Although it was invented in 1993, currently there is no effective cryptanalysis of it found to date. RC4, which means Rivest Cipher 4 or ARC4, was created in 1987 and is a stream cipher. Multiple vulnerabilities have been discovered, thus making it obsolete. The Data Encryption Standard (DES) is a symmetric-key algorithm that has a short key length. Invented in 1975 by IBM, it has had a significant impact on the advancement of cryptography. It has since been replaced by Triple DES, which has also been compromised. Now more advanced encryption standards like Advanced Encryption Standard (AES) are being more commonly deployed. The original name for AES was Rijndael, and it is a symmetric block cipher. AES includes three block ciphers: AES-128, AES-192, and AES-256.

Symmetric is known as a secret key cipher, which uses the same key for encryption and decrypting.

Asymmetric encryption uses two keys to encrypt plaintext. Secret keys are exchanged over the network. This type of encryption ensures that malicious persons do not misuse the keys. It is important to note that anyone with a secret key can decrypt the message, so this is why asymmetric encryption uses two related keys to boost security. A public key is made freely available to anyone who might want to send a message. The second private key is kept a secret so that only the recipient or the sender will know.

A message that is encrypted using a public key can only be decrypted using a private key, whereas a message encrypted using a private key can be decrypted using a public key. Security of the public key is not required because it is publicly available and can be passed over the network or Internet. The asymmetric key has far better power in ensuring the security of information transmitted during communication.

Asymmetric encryption is mostly used in day-to-day communication channels, especially over the Internet. Asymmetric key encryption algorithms include EIGamal, RSA, DSA, elliptic-curve techniques, and PKCS.

Figure 16-10 illustrates asymmetric encryption.

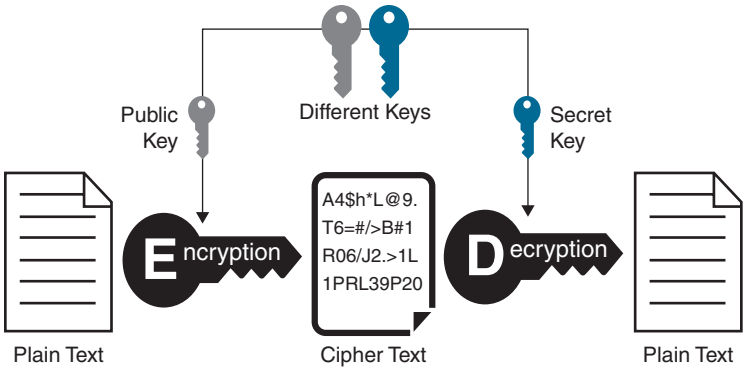


FIGURE 16-10 Asymmetric Encryption

To use asymmetric encryption, you must have a way of discovering public keys. One typical technique is to use digital certificates in a client/server model of communication. A certificate is a package of information that identifies a user and server. It contains information such as an organization's name, the organization that issued the certificate, the user's email address and country, and the user's public key. When a server and client require a secure encrypted communication, they send a query over the network to the other party, which sends back a copy of the certificate. The other party's public key can be extracted from the certificate. A certificate can also be used to uniquely identify the holder.

In an asymmetric key system, each user has a pair of keys: a private key and a public key. To send an encrypted message, you must encrypt the message with the recipient's public key. The recipient then decrypts the message with his or her private key. Remember that public keys encrypt, and private keys decrypt.

Table 16-3 highlights the differences between symmetric and asymmetric encryption.

Table 16-3 Symmetric vs. Asymmetric Encryption

Symmetric Encryption	Asymmetric Encryption
Uses a single key that needs to be shared among the people who need to receive the message, whereas asymmetric encryption uses a pair of public keys and a private key to encrypt and decrypt messages when communicating.	Was introduced to complement the inherent problem of the need to share the key in the symmetric encryption model, eliminating the need to share the key by using a pair of public-private keys.
Is an old technique compared to asymmetric encryption.	Takes relatively more time than symmetric encryption.



Lightweight Cryptography

Lightweight cryptography is an encryption method that features a small footprint and/or low computational complexity. It is aimed at expanding the applications of cryptography to constrained devices such as the ever-expanding Internet of Things (IoT) market. Encryption is an effective countermeasure, and IoT developers are now required to apply encryption to sensor devices in environments with various restrictions that have not previously been subject to encryption. Lightweight cryptography is a technology researched and developed to respond to this issue.

The biggest security-related threat of IoT systems from the traditional IT systems is that even devices used for data collection from the real world can become the target of cyber attacks. For example, the purpose of applying IoT to a manufacturing plant is to significantly improve the productivity and maintainability by collecting data from a large number of sensors installed in production equipment, analyzing it, and performing autonomous control in real time. If sensor data should be falsified during this process, incorrect analysis results would be produced, and erroneous control would result due to such an occurrence having the potential of leading to major damage. Because measurement data and control commands could be trade secrets associated with the know-how of production and management, preventing leakage of information is also important from the viewpoint of competitiveness.

Applying encryption to sensor devices means the implementation of data protection for confidentiality and integrity, which can be an effective countermeasure against threats. Lightweight cryptography enables the application of secure encryption, even for devices with limited resources.

The following factors regarding the implementation are required for lightweight cryptography:

- Size (circuit size, ROM/RAM sizes)
- Power available
- Power consumption
- Processing speed (throughput, delay)

Because the power is greatly dependent on the hardware, such as the circuit size or processor in use, size becomes a reference point for the lightness of the encryption method and for the power. The power consumption depends on the processing speed due to the execution time, so the number of computations that determine the processing speed becomes the index of the lightness. The throughput depends greatly on the parallel processing capability.

Even when the block length and/or secret key length are set shorter than used in standard cryptography, by prioritizing the ease of implementation, such as via a 64-bit block and 80-bit secret key, for example, you still must correctly apply a proven method.

The international standard ISO/IEC 29192, "Lightweight Cryptography," was established at ISO/IEC JTC 1/SC 27. NIST, which issues guidelines on cryptographic technologies, initiated the Lightweight Cryptography Project in 2013 and announced a public call for applications of lightweight cryptographies in 2017.

Key Topic

Steganography

Steganography is the practice of hiding a secret message inside or even on top of something that is not secret. That something can be just about anything you want. These days, many examples of steganography involve embedding a secret piece of text inside a picture or hiding a secret message or script inside a Microsoft Word or Excel document.

The purpose of steganography is to conceal and deceive. It is a form of covert communication and can involve the use of any medium to hide messages. Steganography is not a form of cryptography because it doesn't involve scrambling data or using a key. Instead, it is a form of data hiding and can be executed in unique ways.

Steganography has been used for centuries, but these days, hackers and IT pros have started to use it. The word *steganography* is composed of the Greek root *steganos* for *hidden* or *covered*, and the root *graph* for *to write*. Put these words together, and you have something close to *hidden writing* or *secret writing*. A number of apps can be used for steganography, including Steghide, Foremost, Xiao, Stegais, and Concealment.

TIP The Kali Linux distribution has a number of steganography tools that can be useful in not only hiding data but also detecting and uncovering steganography data.

Audio Steganography

You not only can hide secret information in images and documents but also can hide data and files in audio files. ***Audio steganography*** is a technique used to transmit hidden information by modifying an audio signal in an imperceptible manner. It is the science of hiding some secret text or audio information in a host message.

The host message before steganography and steganography messages after steganography have the same characteristics. Tools like DeepSound allow you to hide and extract secret data or files directly from audio files. Newer tools also support encrypting secret files using AES-256 to improve data protection.

Embedding secret messages in digital sound is a more difficult process. Various techniques for embedding information in digital audio have been established. Audio steganography consists of a carrier or audio file, message, and password. The carrier is also known as a *cover-file*, which conceals the secret information. In the steganography model, the secret message that the sender sends wants it to remain secret. The secret message can be any source: file, text, image, or another audio file. The file is encoded (steg) with the secret key, which is provided to the receiver, who can decode the message and corresponding file. The cover-file with the secret information combined is known as a *stego-file*. Figure 16-11 shows the audio steganography process.

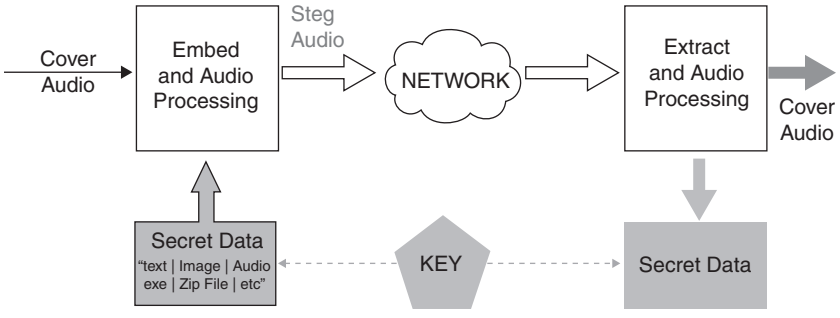


FIGURE 16-11 Audio Steganography Model

Video Steganography

Video steganography is a technique used to hide any kind of file into a cover video file. The use of video-based steganography can be more secure than other multimedia files because of the size and complexity. A video is a collection of frames, and each frame is an image. So if you pull out all the frames from a video, you can use this method to store the data using least significant bit (LSB) steganography and stitch those frames back into a video with the secret message. You can use tools like OpenCv to extract frames from a video, and to extract audio, you could use FFmpeg, which is free and open source. After you have mixed in the steg, you can use FFmpeg to stitch it all together with the hidden message to form a video. How many videos have you watched streamed from the Internet? Now imagine that at least 5 percent of them have some hidden data.

Image Steganography

Image steganography is a technique used to hide any kind of file into an image file. There are currently three effective methods in applying image steganography:

- **Least significant bit (LSB) substitution:** This method works by iterating through the pixels of an image and extracting the Alpha, Red, Green, and Blue (ARGB) hexadecimal values. It then separates the color channels and gets the

least significant bit. Meanwhile, it also iterates through the characters of the message, setting the bit to its corresponding binary value.

- **Blocking:** This method works by breaking up an image into blocks and using Discrete Cosine Transforms (DCT). Each block is broken into 64 DCT coefficients that approximate luminance and color—the values of which are modified for hiding messages.
- **Palette modification:** This method replaces the unused colors within an image's color palette with colors that represent the hidden message.



Homomorphic Encryption

Homomorphic encryption is a form of encryption allowing you to perform calculations on encrypted data without decrypting it first. The result of the computation is in an encrypted form. When decrypted, the output is the same as if the operations had been performed on the unencrypted data.

Common Use Cases

Many use cases combine both symmetric and asymmetric cryptography to improve speed and security at once. Most of these instances use symmetric cryptography to encrypt the bulk of the information and then use asymmetric cryptography to encrypt the symmetric encryption/decryption key (which can in turn be used to decrypt the full message contents).



- **Low-power devices:** Limited by the amount of power or battery available; the most common use cases are mobile devices, tablets, and portable systems. Because of the low-power draw requirements, you can use smaller symmetric key sizes and elliptic-curve asymmetric encryption.
- **Low-latency devices:** Require you to limit the amount of time accessing the CPU, such as instant response time like authentication or read/write. So with that, you need to be able to encrypt and decrypt information rather quickly, and that's where symmetric encryption and smaller keys come in. Some block ciphers like PRINCE that use a 64-bit block cipher with a 128-bit key perform well.
- **High-resiliency applications:** Generally require an encryption method that provides integrity of the data being exchanged. You would want to use large keys and a strong encryption method that would also include hashing to ensure that integrity.

- **Supporting confidentiality:** One of the main reasons for using encryption is to make sure your data remains confidential, and you have some guarantee of secrecy and privacy. It's common to use drive-level, file-level, and even email encryption.
- **Supporting integrity:** When you are sending information to someone else, you want to make sure the information that person receives is the original information you sent and wasn't modified. Implementation of this feature guarantees that information wasn't tampered with. It's common to use hashes to provide the integrity, where you take the hash prior to sending and then verify on the receiving end. You use integrity to prevent modification of data, protect file downloads. and store passwords.
- **Supporting obfuscation:** This approach can be used with cryptography to obfuscate or hide data. Modern malware has been used to embed secret messages in state and corporate espionage. Enterprises often need to hide data from unauthorized viewing and access, especially business-critical data or personal information. Obfuscation can be used for data security reasons or compliance related to data protection. Where there is a need for a full data set, obfuscation is used to preserve privacy. This need might exist for several reasons, such as data exports and secure transactions.
- **Supporting authentication:** This approach is always used with authentication, where it takes the original password and takes a hash of it. In more modern methods, it also adds a salt to help randomize the stored password hashes. In this case, if someone gains access to the stored password hashes on the system, the salt ensures identical passwords are still different, making it harder to crack.
- **Supporting nonrepudiation:** This approach allows you to confirm or prove the authenticity of any information you received from a third party—that it did really come from the sender. Using digital signatures can provide integrity and nonrepudiation of data that has been sent/received.

Limitations

The limitations of cryptography in action are typically specific to a platform, such as IoT devices, with low power and limited CPU capabilities, meaning the key size may need to be smaller, be replaced more often, and take longer to encrypt data in/out of the device. Using weak keys or implementations of specific cryptography itself might be more acceptable if there are expected off-device protections (such as firewalls and other controls). Limitations to keep in mind include the following:



- **Speed:** Encryption is taxing on a system, so if speed of data encryption and decryption is important, you need higher power, a more capable CPU, and more memory. The trade-off comes when you have smaller embedded systems with limited power, smaller CPUs, and less memory.

- **Size:** Key size matters. When used properly and to its full extent, encryption requires large keys and higher levels of encryption. As we have seen, a larger key with the right algorithm and cipher provides much higher-level protection, whereas having the wrong cipher and a large key can lead to guessable keys.
- **Weak keys:** These are the same as weak passwords: they allow a system to be easily compromised. Weak keys are generally the result of poorly implemented cryptography, not in the standard but in the manner in which it has been implemented in a particular technology.
- **Time:** As related to cryptography, time is extremely important because timing attacks use side channels to analyze the time taken to execute a cryptographic algorithm. Time can also be used in generating keys used in cryptography; these have been mostly replaced with more randomness.
- **Longevity:** This factor concerns the lifecycle of a cryptographic key. After a key is generated, the key management system should control the sequence of states that a key progresses over its lifecycle and allow an authorized administrator to manage them. From NIST, a crypto period is the operational life of a key and is determined by a number of factors, such as the sensitivity of the data or keys to be protected, and how much data or how many keys are being protected.
- **Predictability:** Predictability is the enemy of secure cryptography because the entire cryptographic system is only as good or strong as the source of randomness used to generate the primes. Prime numbers are whole numbers greater than 1 and are divisible only by the number 1 or itself. Random number generators come in two basic forms: software and hardware. Software solutions are not capable of providing true randomness because they are based on deterministic computer programs, such as time. Most hardware random number generators rely on classical physics to produce what looks like a random stream of bits. However, in reality, determinism is hidden behind complexity.
- **Reuse:** This is a characteristic of a generally badly developed algorithm. Reuse of keys is considered a bad practice because of the potential for exact results, collisions and similar data, keys, and decryption producing the same results, leading to weak keys.
- ***Entropy*:** This is the foundation upon which all cryptographic functions operate. Entropy, in cybersecurity, is a measure of the randomness or diversity of a data-generating function. Data with full entropy is completely random, and no meaningful patterns can be found.
- **Computational overheads:** Current constructions of cryptographic primitives involve a large multiplicative computational overhead that grows with the desired level of security. There is constant growth when implementing higher-level security with cryptographic primitives with a constant computational overhead.

- **Resource vs. security constraints:** The more secure the encryption used and the higher the key length, the more processing power and memory that the server will need. The goal in any implementation should be to strike a balance between the hardware resources that the server has and the amount of processing power.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 16-4 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 16-4 Key Topics for Chapter 16

Key Topic Element	Description	Page number
Section	Digital Signatures	395
Section	Key Length	396
Section	Key Stretching	397
Section	Salting	397
Section	Hashing	398
Section	Key Exchange	399
Section	Elliptic-Curve Cryptography	399
Section	Perfect Forward Secrecy	400
Section	Quantum	401
Section	Post-Quantum	402
Section	Ephemeral	403
Section	Modes of Operation	403
Section	Counter Mode	408
Section	Blockchain	409
Section	Cipher Suites	410

Key Topic Element	Description	Page number
Section	Symmetric vs. Asymmetric Encryption	411
Section	Lightweight Cryptography	414
Section	Steganography	415
Section	Homomorphic Encryption	417
List	How to apply encryption for common use cases	417
List	Various factors that affect the functionality of cryptography	418

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

digital signatures, key length, key stretching, salting, hash, key exchange, elliptic-curve cryptography (ECC), perfect forward secrecy, quantum cryptography, post-quantum cryptography, ephemeral keys, authenticated mode, unauthenticated mode, counter mode, blockchain, public ledger, cipher suite, stream cipher, block ciphers, symmetric encryption, asymmetric encryption, lightweight cryptography, steganography, audio steganography, video steganography, image steganography, homomorphic encryption, entropy

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. Digital signatures employ which stream type?
2. What are key stretching techniques used for?
3. What does salting passwords protect against?
4. What type of key do block ciphers use?
5. What type of encryption is known as public key cryptography?
6. With an asymmetric key system, to send an encrypted message to someone, what must you encrypt the message with?
7. What kind of key is designed to be used for a single transaction or session?



This chapter covers the following topics related to Objective 3.1 (Given a scenario, implement secure protocols) of the CompTIA Security+ SY0-601 certification exam:

- Protocols
  - Domain Name System Security Extensions (DNSSEC)
  - SSH
  - Secure/Multipurpose Internet Mail Extensions (S/MIME)
  - Secure Real-Time Transport Protocol (SRTP)
  - Lightweight Directory Access Protocol over SSL (LDAPS)
  - File Transfer Protocol, Secure (FTPS)
  - SSH File Transfer Protocol (SFTP)
  - Simple Network Management Protocol, version 3 (SNMPv3)
  - Hypertext Transfer Protocol over SSL/TLS (HTTPS)
  - IPsec
    - Authentication header (AH)/Encapsulating Security Payloads (ESP)
    - Tunnel/transport
  - Post Office Protocol (POP)/ Internet Message Access Protocol (IMAP)
- Use Cases
  - Voice and Video
  - Time Synchronization
  - Email and web
  - File transfer
  - Directory services
  - Remote access
  - Domain name resolution
  - Routing and switching
  - Network address allocation
  - Subscription services

Implementing Secure Protocols

A security protocol is an abstract of a protocol that performs a function in a secure manner, one that applies cryptographic methods, often as sequences of cryptographic primitives. Using a secure protocol such as TLS is essential when connecting to a password-protected URL. If the connection is done through a plaintext protocol, eavesdroppers could obtain passwords with little effort. In particular, all levels of network protocols are completely exposed to eavesdropping and manipulation by attackers, so using cryptography properly is a requirement to derive any benefit.

Effective security mechanisms protect session-oriented and session establishment protocols. Although they are different, they can share many cryptographic primitives. Many insecure protocols have secure equivalents, as we discuss at greater length in this chapter. As you learn here, you can often find a secure protocol with the type of traffic you wish to carry. Instead of operating over the command line with Telnet, you can use Secure Shell (SSH), and instead of transferring files with FTP, you can use Secure File Transfer Protocol (SFTP), which is also based on SSH.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 17-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 17-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Protocols	1–7
Use Cases	8–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. How would you create a chain of trust when using DNS?
   a. Configure a CNAME record and point to MS5.
   b. Set up SPF records in DNS with a metric of 5.
   c. Configure and deploy SDNS.
   d. Configure and use DNSSEC.

2. In which way would you be able to secure LDAP?
   a. Install Microsoft LDAP on all AD Primary Domain Controllers.
   b. Install a properly formatted PKI certificate.
   c. Uninstall LDAP and install secure LDAP.
   d. Use a site connector in AD and select Secure.

3. How does using SSH help secure communications?
   a. By automatically disabling TFTP and Telnet
   b. By using asymmetric (public key) cryptography
   c. By using an encryption key fob and special code
   d. By configuring a shell on the remote host

4. S/MIME utilizes which cryptographic mathematic key type?
   a. Symmetric
   b. Asymmetric
   c. Biometric Sequence
   d. AD Key Encryption

5. Authentication Header (AH)/Encapsulating Security Payloads (ESP) optional packet header is used to guarantee which of the following?
   a. Packet delivery in congested networks
   b. QoS in cryptographic traffic
   c. Connectionless integrity
   d. Data origin packet forwarding

6. Which email protocol allows the user to download all the messages and still retain a copy of the email on the email server?
   a. POP3
   b. NMAP
   c. IMAPS
   d. ADFS

7. What Hypertext Transfer Protocol provides encrypted communications?
   a. SMTP
   b. NMAP
   c. IMAPS
   d. HTTPS

8. Time synchronization begins with configuring NTP on servers, devices, and hosts. Which time zone should you configure your NTP server to utilize?
   a. CST/CDT (Central time zone)
   b. EST/EDT (Eastern time zone)
   c. PST/PDT (Pacific time zone)
   d. UTC (Coordinated Universal Time)

9. What type of remote access is the most widely used and secure method considering all your assets are located behind a firewall?
   a. VPN
   b. APN
   c. Private line terminal
   d. A rate-limited policy applied to inbound interfaces

10. A Domain Name System (DNS) has a secure counter protocol called DNSSEC. If you can't implement DNSSEC, which of the following are ways to secure DNS without migrating to DNSSEC? (Select all that apply.)
    a. Limit or disable zone transfers.
    b. Use trusted DNS forwarders.
    c. Use firewalls and limit communications to the DNS server.
    d. Use caching-only DNS servers.

Foundation Topics

Protocols

A protocol is a standard set of rules that allow electronic devices to communicate with each other. These rules include what type of data may be transmitted, what commands are used to send and receive data, and how data transfers are confirmed. You can think of a protocol as a spoken language. If two or more people share the same language, they can communicate effectively. Similarly, if two hardware devices support the same protocol, they can communicate with each other, regardless of the manufacturer or type of device.

Key Topic

Domain Name System Security Extensions

The job of the Domain Name System (DNS) is to translate IP addresses into names via UDP port 53. The lack of security in DNS has been a well-known and documented issue for a decade, one that the security community has been trying to solve since the late 1990s. These efforts resulted in the creation of ***DNS Security Extensions (DNSSEC)***. DNSSEC provides secure answer validation. It does this through public key cryptography; it is backward-compatible and can be deployed side by side with traditional DNS. DNSSEC does not encrypt traffic. Although DNSSEC uses both public and private keys, they are used only for authentication purposes. In addition to being able to stop cache poisoning and other DNS-based attacks, DNSSEC also enhances other security features. Many organizations already publish antispam information such as Sender Policy Framework (SPF) or Domain Keys Identified Mail (DKIM) in DNS.

DNSSEC is used in securing the chain of trust that exists between the DNS records that are stored at each domain level, verifying each trust between the child level and its parent, all the way back to the root zone. Through this multilevel process, the integrity of the DNS records associated with a domain can be verified, thus ensuring to the client that the website or service requested and the one delivered are, in fact, one and the same.

Accessing validated websites and services is the aim of security extension-enabled DNS services. The goal here is to reach the intended servers hosting those sites or services. As far as protection goes, it ensures against malicious URLs designed to impersonate a site or service for the purpose of harvesting account names and passwords. This could come in the form of a maliciously injected record during an on-path/man-in-the-middle attack or as part of a known vulnerability, such as DNS-cache poisoning or spoofing attacks. In either case, DNSSEC will reply with a 404 error (website not found) in the event that a domain does not resolve due to DNS records that can't be validated.

To set up DNSSEC, follow these steps:

Step 1. Verify that your top level domain (TLD) supports the DNS security extensions.

Step 2. Work with your service provider to ensure it will support DNSSEC.

Step 3. Generate the zone signing key (ZSK) and key signing key (KSK) for your domain DNS zone.

Step 4. Sign your DNS zone to generate signed zone records for your domain(s).

Step 5. Generate the DS Declaration of Signing Record, which contains hashed values for the cryptographic keys used to sign your DNS zone.

Step 6. Import the DS record(s) for your domain(s) to the self-hosted or fully managed nameserver, ensuring that the information is obtained in the second step.

Step 7. Test your chain of trust and DNSSEC.

NOTE DNS was designed in the 1980s when the Internet and the services it offered were in their infancy and security was not a primary consideration in the design.

TIP DNSSEC is defined in RFCs 4470, 4641, 5155, and 6014.

Key Topic

SSH

Secure Shell (SSH) is a network communication protocol that allows two computers or devices to communicate and share data. An inherent feature of SSH is that the communication between the two systems is encrypted, making it suitable for use on insecure networks. SSH replaces Telnet, an insecure protocol that was used to connect to systems and devices. SSH uses asymmetric (public key) RSA cryptography for both connection and authentication.

Implementing SSH across your organization should be a phased approach, starting with high-value/high-profile assets first, then systems, and finally networking devices. Figure 17-1 illustrates the SSH connection process.

There are three distinct parts needed to utilize SSH: the **ssh** command; the user, such as root; and the host, a hostname, or an IP address:

```
ssh {user}@{host}
```

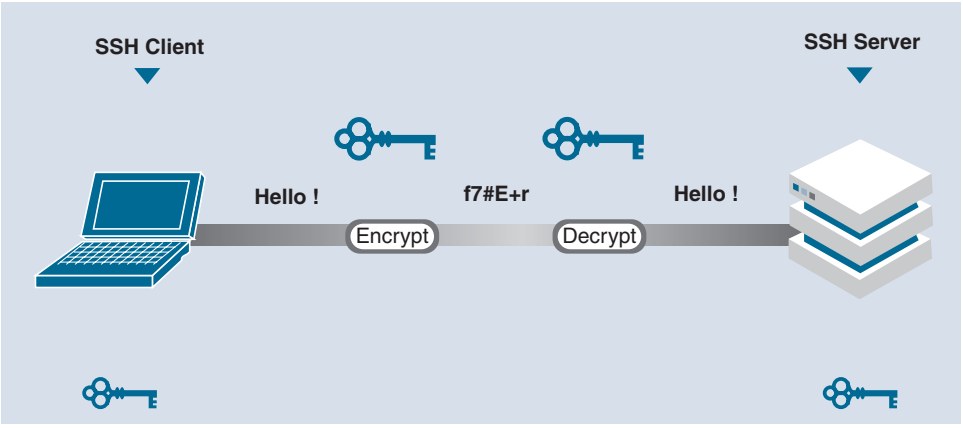


FIGURE 17-1 The SSH Connection Process

The SSH key command instructs your system that you want to open an encrypted Secure Shell connection. {*user*} represents the account you want to access. For example, you may want to access the root user, which is basically synonymous with system administrator with complete rights to modify anything on the system. {*host*} refers to the computer you want to access. It can be an IP address (e.g., 8.6.75.30) or a domain name (e.g., www.xyzdomain.com).

After you press Enter, you are prompted to enter the password for the requested account. When you type it in, nothing appears on the screen, but your password is, in fact, being transmitted. When you're done typing, press Enter once again. If your password is correct, you are greeted with a remote terminal window.

The significant advantage offered by SSH over its predecessors is the use of encryption to ensure secure transfer of information between the host and the client. Host refers to the remote server you are trying to access, while the client is the computer you are using to access the host. SSH uses three different encryption technologies:

- Symmetrical encryption
- Asymmetrical encryption
- Hashing

Key Topic

Secure/Multipurpose Internet Mail Extensions

Secure/Multipurpose Internet Mail Extensions (S/MIME) is a technology that enables you to encrypt your emails. S/MIME is based on asymmetric cryptography to protect your emails from snoopers and unwanted access. It also allows you to digitally sign your emails to verify you as the legitimate sender of the message,

making it an effective weapon against many phishing attacks out there. That's basically the core of what S/MIME is all about. S/MIME is based on asymmetric cryptography, which uses a pair of mathematically related keys to operate: a public key and a private key. Emails are encrypted with the recipient's public key. The email can be decrypted only with the corresponding private key, which is supposed to be in the sole possession of the recipient. Unless the private key is compromised, you can be confident that only your intended recipient will be able to access the sensitive data in your emails.

Figure 17-2 illustrates how S/MIME certificates utilize asymmetric encryption to encrypt and decrypt plaintext information, where the sender uses the public key of the recipient to send a message.

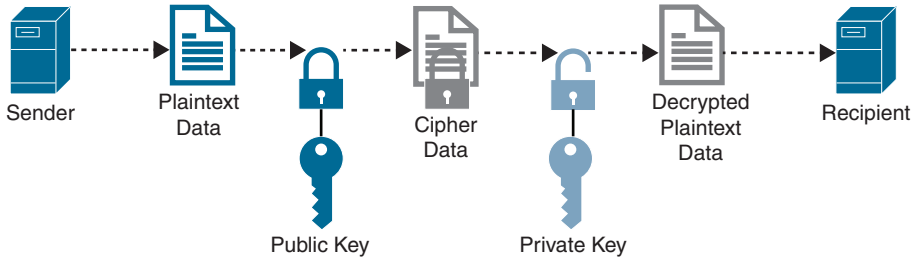


FIGURE 17-2 S/MIME Certificate Process

Let's look at how asymmetric encryption works by using this example. In this scenario Chris wants to send an email to his friend Joseph but does not want anyone else on the network to read or modify it. Chris uses Joseph's public key to encrypt the message. Anyone can have access to this public key. Therefore, Chris or anyone can encrypt the message they want to send to Joseph. An intruder/attacker on the network cannot read this message, though, without the corresponding private key. The private key is accessible only to Joseph, so only he can decrypt and read the message. You can get an email signing certificate from a trusted certificate authority (CA) and use it to digitally sign your messages (the certificate is automatically attached to the mail) before sending it off. This process works as follows:

1. A digital signature is associated with two keys: a private key and a public key.
2. Authentication is done using the public key, whereas the private key is used to generate the signature itself.
3. The public key is sent along with every protected email message to assert the identity of the sender.
4. The private key, on the other hand, generates and applies the unique digital signature to each email. The signature verifies that the message is unaltered and that an unauthorized third party has not tampered with it.



Secure Real-Time Transport Protocol

Secure Real-time Transport Protocol (SRTP) is the secure version of Real-time Transport Protocol (RTP) intended to provide encryption, message authentication, and integrity, and to provide replay attack protection to the RTP data in both unicast and multicast applications. SRTP and SRTCP use the Advanced Encryption Standard (AES) as the default cipher. SRTP was initially published by the Internet Engineering Task Force (IETF) in RFC 3711 in March 2004.

SRTP provides confidentiality by encrypting the RTP payload, not including the RTP header. It supports source origin authentication and is widely used as the security mechanism for RTP. Although it can be used in its entirety, specific security features can be disabled. The main caveat with SRTP is key management. There are many options, including DTLS-SRTP, MIKEY in SIP, Security Description (SDES) in SDP, ZRTP, and others.

The two defined cipher modes allow the AES block cipher to be used as a stream cipher:

- **Segmented Integer Counter Mode:** In general, almost any function can be used in the role of counter, assuming that this function does not repeat for a large number of iterations. But the standard for encryption of RTP data is just the usual integer incremental counter. AES running in this mode is the default encryption algorithm, with a default key size of 128 bits and a default session salt key length of 112 bits.
- **f8-mode:** This is a variation of output feedback mode, enhanced to be seekable and with an altered initialization function. The default values of the encryption key and salt key are the same as for AES in counter mode. Besides the AES cipher, SRTP enables you to disable encryption outright, using the NULL cipher, which does not perform any encryption. The encryption algorithm functions as the identity function and copies the input stream to the output stream without any changes. It is mandatory for this cipher mode to be implemented in any SRTP-compatible system.

Figure 17-3 illustrates the SRTP process, which adds further security features, such as message authentication, confidentiality, and replay protection.

Though SRTP can easily accommodate new encryption algorithms, the SRTP standard states that new encryption algorithms may only be introduced through publication of a new companion standard-track RFC, which must clearly define the new algorithm. Figure 17-4 illustrates the SRTP packet structure compared to the standard RTP packet. SRTP adds the Auth and MKI master key identifiers to the packet.

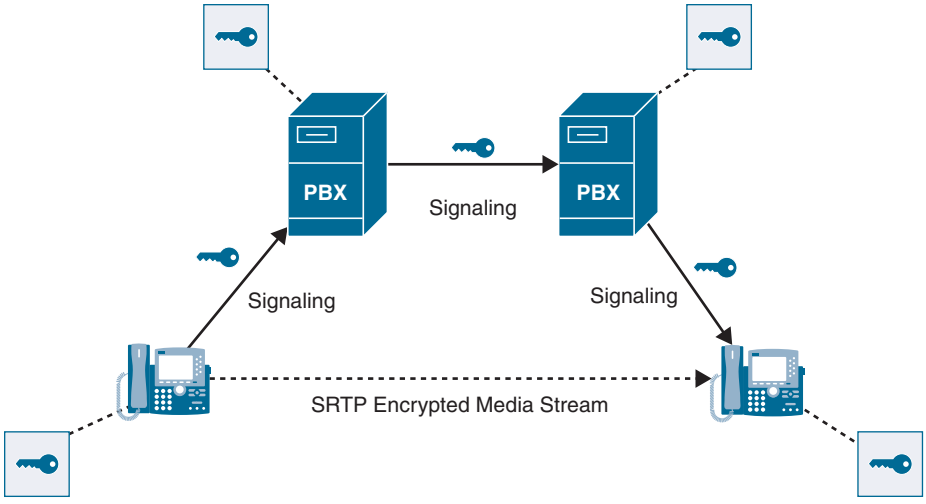


FIGURE 17-3 The SRTP Process

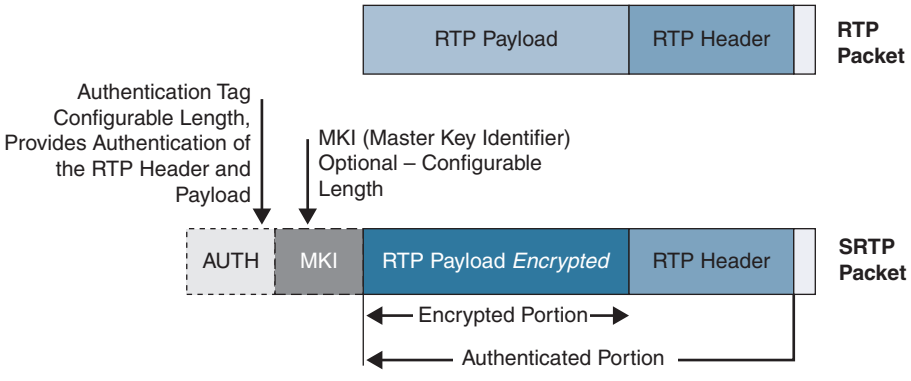


FIGURE 17-4 The SRTP Packet Structure with AUTH and MKI

The MKI is optional and indicates the master key ID, which is used to generate SRTP session keys. It is useful in cases that support keying, changing the key during the lifetime of the session. The MKI points to the ID of the master key to be used, while the actual key is delivered elsewhere by using a key management mechanism.

NOTE SRTP is normally used to secure voice and video transmissions. Voice and video calls are established with Session Initiation Protocol (SIP), and data is transmitted with RTP.

TIP SRTP is based on RFCs 2104, 3711, 4771, and 3550 and falls under cryptographic protocols.

Key Topic

Lightweight Directory Access Protocol over SSL

LDAP is used to read and write information to Active Directory. By default, LDAP traffic is transmitted unsecured. You can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL)/Transport Layer Security (TLS) technology. You can enable ***Lightweight Directory Access Protocol over SSL (LDAPS)*** by installing a properly formatted certificate from a certificate authority (CA) according to the guidelines. LDAPS over SSL/TLS uses TCP port 636.

There's no user interface for configuring LDAPS. Installing a valid certificate on a domain controller permits the LDAP service to listen for and automatically accept SSL connections for both LDAP and global catalog traffic.

To enable LDAPS, you must install a certificate that meets the following requirements:

- The LDAPS certificate is located in the local computer's personal certificate store (known as the computer's MY certificate store).
- A private key that matches the certificate is present in the local computer's store and is correctly associated with the certificate.
- The Enhanced Key Usage extension includes the Server Authentication (1.3.6.1.5.5.7.3.1) object identifier (also known as OID).
- The Active Directory fully qualified domain name (FQDN) of the domain controller must appear in one of the following places:
  - The Common Name (CN) in the Subject field.
  - A DNS entry in the Subject Alternative Name extension.
- The certificate was issued by a CA that the domain controller and the LDAPS clients trust. Trust is established by configuring the clients and server to trust the root CA to which the issuing CA chains.
- You can use the Schannel cryptographic service provider (CSP) to generate the key.

Key Topic

File Transfer Protocol, Secure

File Transfer Protocol, Secure (FTPS) is also known as FTP over SSL, FTP/SSL, FTP-SSL, FTP-ES, and FTP Secure. FTPS is a name used to encompass a number

of ways in which FTP software can perform secure file transfers. Each involves the use of a security layer below the standard FTP protocol to encrypt data. FTPS uses multiple port numbers for implicit and explicit connection types, so every time a file transfer or directory listing request is made, another port opens. When file transfers are sent, they are exchanged using FTPS and can be authenticated through FTPS-supported methods like passwords, client certificates, and server certificates.

Negotiation is not supported with implicit FTPS configurations. A client is immediately expected to challenge the FTPS server with a TLS Client Hello message. If the FTPS server does not receive such a message, the server should drop the connection.

To maintain compatibility with existing non-FTPS-aware clients, implicit FTPS was expected to listen on the IANA well-known port 990/TCP for the FTPS control channel and port 989/TCP for the FTPS data channel. This way, administrators were able to retain legacy-compatible services on the original 21/TCP FTP control channel.

NOTE That implicit negotiation was not defined in RFC 4217. As such, it is considered an earlier, deprecated method of negotiating TLS/SSL for FTP.

In explicit mode, FTPES, an FTPS client, must "explicitly request" security from an FTPS server and then step up to a mutually agreed encryption method. If a client does not request security, the FTPS server can either allow the client to continue in insecure mode or refuse the connection.

The mechanism for negotiating authentication and security with FTP was added under RFC 2228, which included the new FTP command **AUTH**. Although this RFC does not explicitly define any required security mechanisms, SSL or TLS, it does require the FTPS client to challenge the FTPS server with a mutually known mechanism. If the FTPS client challenges the FTPS server with an unknown security mechanism, the FTPS server responds to the AUTH command with error code 504 (not supported). Clients may determine which mechanisms are supported by querying the FTPS server with the **FEAT** command, although servers are not necessarily required to be honest in disclosing what levels of security they support. Common methods of invoking FTPS security included AUTH TLS and AUTH SSL.

NOTE The explicit method is defined in RFC 4217. In later versions of the document, FTPS compliance required that clients always negotiate using the AUTH TLS method.

Secure (or SSH) File Transfer Protocol

Secure File Transfer Protocol (SFTP) is a file protocol for transferring large files over the web. This term is also known as *Secure Shell (SSH) File Transfer Protocol*. It builds on the File Transfer Protocol (FTP) and includes Secure Shell (SSH) security components.

Secure Shell is a cryptographic component of Internet security. SSH and SFTP were designed by the Internet Engineering Task Force for greater web security.

There are two different ways to use SFTP as a protocol. The first one is as a graphical user interface, or GUI, where a program abstracts the use of SFTP visually for end users. Using this method, you can often drag and drop files or use control buttons to send files with this protocol over the Internet. Administrators who are designing systems for end users typically use a GUI method to help others achieve SFTP protocol use goals.

The second way to use SFTP is at the command line. A user would utilize a command-line command for SFTP. This approach is often used in a Linux environment. With this type of setup, the user has to type in specific command lines to generate the SFTP protocol.

SFTP, as a successor to FTP, is used for many situations where file security is important. One of the biggest ones is to comply with standards like the federal Health Insurance Portability and Accountability Act (HIPAA) that governs protected health information (PHI). Any business, even a third party working with a hospital or health-care provider, must keep its PHI confidential, and that includes during its transition through networks in digital packet form. That's why SFTP can be useful in securing this type of data. SFTP is one of several options for shielding that data in transfer, to make sure that hackers don't obtain it, and that the company does not unwittingly perform a HIPAA violation if that law applies.

SFTP can also satisfy other standards for data protection. For example, the General Data Protection Regulation (GDPR) governing the data of European Union (EU) users may require this type of security. SFTP can also be handy in other situations where sensitive data needs to be protected. For example, trade secrets might not be covered by any particular data privacy rule, but it can be devastating for them to fall into the wrong hands. So a business user might use SFTP to transmit files containing trade secrets or other similar information.

Simple Network Management Protocol Version 3

The ***Simple Network Management Protocol version 3 (SNMPv3)*** feature provides secure access to devices by authenticating and encrypting data packets over the network. SNMPv3 is an interoperable, standards-based protocol that is defined in RFCs 3413 to 3415.

Security features provided in SNMPv3 are as follows:

- **Message integrity:** Ensures that a packet has not been tampered with during transit
- **Authentication:** Determines that the message is from a valid source
- **Encryption:** Scrambles the content of a packet to prevent it from being learned by an unauthorized source

SNMPv3 is a security model in which an authentication strategy is set up for a user and the group in which the user resides. Security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is used when handling an SNMP packet. The following configuration types are available when configuring SNMPv3:

- **No authentication and no privacy (noAuthNoPriv):** Usually used for monitoring
- **Authentication and no privacy (authNoPriv):** Usually used for control
- **Authentication and privacy (authPriv):** Usually used for downloading secrets

To configure SNMPv3 on most Cisco devices, you first enter **snmp-server group group-name v3 auth**. Configuration on other devices normally starts with **snmp.**

Example 17-1 shows the configuration of Cisco IOS-XE and Cisco ASA firewall samples.

Example 17-1 Configuring SNMPv3

```
Cisco IOS, IOS-XE, IOS-XR and OS-NX
snmp-server community <your-community> RO

Cisco ASA
snmp-server community <your-community> snmp-server contact
<your-contact>
snmp-server location <your-location> snmp-server host <interface>
<NMS-IP>
 poll community <your-community> version 2c
```

To verify your devices have been properly configured, you can use an open-source tool called SNMPWALK.

Example 17-2 shows the **snmpwalk -v3** command help and options available when using the tool.

Example 17-2 snmpwalk -v3 Command Help

```
snmpwalk -v3 -l <noAuthNoPriv|authNoPriv|authPriv> -u <username>
[-a <MD5|SHA>] [-A <authphrase>] [-x <DES|AES>] [-X <privaphrase>]
<ipaddress>[:<dest_port>] [oid]


# snmpwalk v3 example with authentication and encyrption
snmpwalk -v3 -l authPriv -u UserMe -a SHA -A AuthPass1 -x AES
-X PrivPass2 192.168.10.1:161 1.3.6.1.2.1.1
```

To install SNMPWALK on most Debian-based operating systems, you install Net-SNMP by using the **apt-get install snmp snmpwalk** command.

SNMP uses what is known as Management Information Bases (MIBs) to locate specific elements of reporting the status on a device. An MIB is a collection of information organized hierarchically that is accessed using a protocol such as SNMP. MIBs are collections of definitions that define the properties of the managed objects within the device to be managed.

Key Topic

Hypertext Transfer Protocol over SSL/TLS

Hypertext Transfer Protocol Secure (HTTPS) is the secure version of HTTP, which is the primary protocol used to send data between a web browser and a website. HTTPS is encrypted to increase security of data transfer. This encryption is particularly important when users transmit sensitive data, such as by logging in to a bank account, email service, or health insurance provider.

Any website, especially one that requires login credentials, should use HTTPS. In modern web browsers such as Chrome, websites that do not use HTTPS are marked differently from those that are. Look for a green padlock in the URL address bar to signify the web page is secure. Web browsers take HTTPS seriously and show sites that are insecure with a blocked lock, and now both Chrome and Firefox do not allow you to see an unsecure page unless you agree to a "lack of security" notice.

To encrypt communications, HTTPS uses an encryption protocol called Transport Layer Security (TLS), although formerly it was known as Secure Sockets Layer (SSL). This protocol secures communications by using what's known as an asymmetric public key infrastructure. This type of security system uses two different keys to encrypt communications between two parties:

- **The private key:** This key is controlled by the owner of a website, and it's kept private and secure away from potential attackers. This key lives on a web server and is used to decrypt information encrypted by the public key.

- **The public key:** This key is available to everyone who wants to interact with the server in a way that's secure. Information that's encrypted by the public key can be decrypted only by the private key.

HTTPS prevents websites from having their information broadcast in a way that's easily viewed by anyone snooping on the network. HTTPS uses port 443 for secure communications. When information is sent over regular HTTP, port 80, it is unencrypted; the information is broken into packets of data that can be easily "sniffed" using free software such as Wireshark. This makes communication over the unsecure medium, such as public Wi-Fi, highly vulnerable to interception. All communication that occurs over HTTP occurs in plaintext, making it highly accessible to anyone with the correct tools, and vulnerable to on-path attacks.



IPsec

IPsec, also known as Internet Protocol Security or IP Security protocol, defines the architecture for security services for IP network traffic. IPsec describes the framework for providing security at the IP layer, as well as the suite of protocols designed to provide that security, through authentication and encryption of IP network packets. Also included in IPsec are protocols that define the cryptographic algorithms used to encrypt, decrypt, and authenticate packets, as well as the protocols needed for secure key exchange and key management.

IPsec originally defined two mechanisms for imposing security on IP packets: the Encapsulating Security Payload (ESP) protocol, which defined a method for encrypting data in IP packets, and the Authentication Header (AH) protocol, which defined a method for digitally signing IP packets. The Internet Key Exchange (IKE) protocol is used to manage the cryptographic keys used by hosts for IPsec.

Authentication Header/Encapsulating Security Payloads

IPsec uses two distinct protocols: AH and ESP. The AH protocol provides a mechanism for authentication only. AH provides data integrity, data origin authentication, and an optional replay protection service.

- The IP ***Authentication Header (AH)***, specified in RFC 4302, defines an optional packet header to be used to guarantee connectionless integrity and data origin authentication for IP packets, and to protect against replays.
- The IP ***Encapsulating Security Payload (ESP)***, specified in RFC 4303, defines an optional packet header that can be used to provide confidentiality through encryption of the packet, as well as integrity protection, data origin authentication, access control, and optional protection against replays or traffic analysis.

Tunnel/Transport

IPsec can be used to protect network data, for example, by setting up circuits using IPsec tunneling, in which all data being sent between two endpoints is encrypted, as with a virtual private network (VPN) connection; for encrypting application layer data; and for providing security for routers sending routing data across the public Internet. IPsec can also be used to provide authentication without encryption—for example, to authenticate that data originates from a known sender. ESP and AH can operate between hosts and even between networks. They can also operate in two modes: the less-secure Transport mode that encrypts the data packet, for use between two workstations that are running a VPN client; and Tunnel mode, which is more secure, encrypts the whole packet including header info and source, and is used between networks.

- ***Tunnel mode*:** Usually used between secured network gateways, IPsec Tunnel mode enables hosts behind one of the gateways to communicate securely with hosts behind the other gateway. For example, any users of systems in an enterprise branch office can securely connect with any systems in the main office if the branch office and main office have secure gateways to act as IPsec proxies for hosts within the respective offices. The IPsec tunnel is established between the two gateway hosts, but the tunnel itself can carry traffic from any hosts inside the protected networks. Tunnel mode is useful for setting up a mechanism for protecting all traffic between two networks, from disparate hosts on either end.
- ***Transport mode*:** When two individual hosts set up a directly connected IPsec VPN connection, the circuit can be said to be an example of a Transport mode IPsec circuit. For example, a Transport mode IPsec circuit might be set up to allow a remote IT support technician to log in to a remote server to do maintenance work. Transport mode IPsec is used in cases where one host needs to interact with another host; the two hosts negotiate the IPsec circuit directly with each other, and the circuit is usually torn down after the session is complete.



Post Office Protocol/Internet Message Access Protocol

Internet Message Access Protocol (IMAP) allows you to access your email wherever you are, from any device. IMAP has been around for a long time; the current version is IMAPv4 (IMAP4). When you read an email message using IMAP, you aren't actually downloading or storing it on your computer; instead, you're reading it from the email service. As a result, you can check your email from different devices, anywhere in the world: your phone, a computer, a friend's computer. IMAP downloads a message only when you click on it, and attachments aren't automatically downloaded.

This way, you're able to check your messages a lot more quickly than when using ***Post Office Protocol (POP)***. IMAP operates on port 143 nonencrypted and port 993 SSL/TLS known as IMAPS.

POP is in its third revision, called POP3, and it works by contacting your email service and downloading all of your new messages from it. After they are downloaded onto your PC, they are deleted from the email service. This means that after the email is downloaded, you can only access it using the same computer. Sent mail is stored locally on your system, not on the email server. POP3 operates on TCP port 110 nonencrypted and port 995 SSL/TLS, known as POP3S.

Figure 17-5 illustrates how an email, using both secure and insecure protocols and ports, would flow from a client sender to a client receiver. There are many options when configuring your client for receiving and sending emails; however, only the secure methods are acceptable practice nowadays.

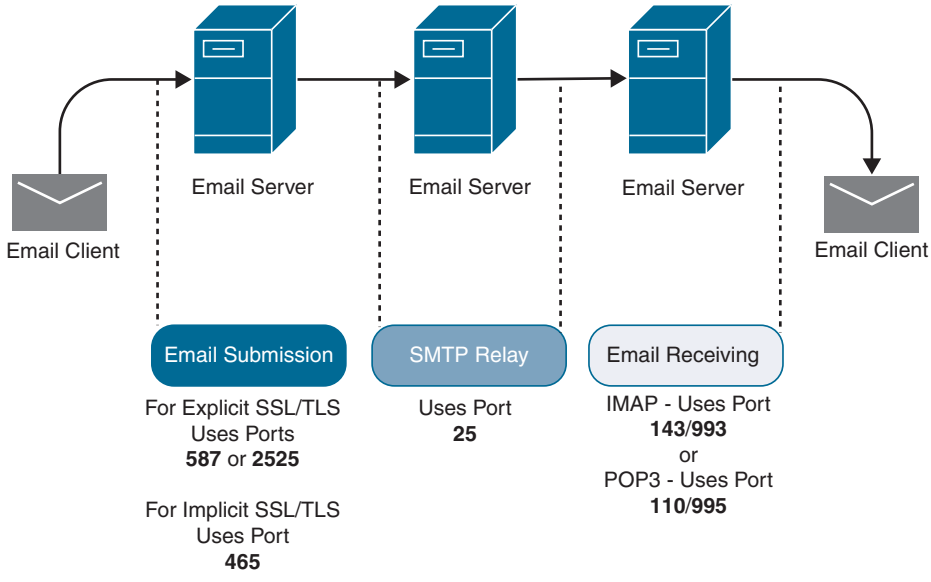


FIGURE 17-5 Email Flow and Security Options Between Sender and Receiver

Both POP and IMAP are insecure protocols by default; the good news is that they have secure port alternatives, as mentioned previously.

Key Topic

Use Cases

Protocols are a defined set of rules that allow different components to have a "common language" to exchange commands and data. Having defined protocols

allows for the development of interoperable devices. Secure protocols have built-in security mechanisms.

Voice and Video

Deploying secure voice and video is essential to ensuring your communications are private. To properly secure your voice and video networks, you should set up secure private VLANs; configure quality of service (QoS), anti-brute-force technology, and two-factor authentication (where available); and configure and deploy firewalls to secure these systems and networks. SRTP is ideal for protecting Voice over IP (VoIP) traffic because it can be used in conjunction with header compression and has no effect on IP quality of service. This provides significant advantages, especially for voice traffic using low-bitrate voice codecs such as G.729. SRTP can also be used with ITU Recommendation H.264 and MPEG-4 to stream video securely in multimedia applications.

Time Synchronization

The Network Time Protocol (NTP) synchronizes the time of a computer client or server to another server or within a few milliseconds of Coordinated Universal Time (UTC). NTP servers, long considered a foundational service of the Internet, are used on nearly every computer, whereas the lack of time synchronization on a network creates an opportunity for replay attacks, which involve a fake or malevolent repeated delay of an authentic data transmission. NTP uses UDP to synchronize time based on an atomic clock. NTP servers use UDP port 123, and NTP clients use random ports above 1023.

Enterprises looking to deploy time synchronization should utilize three public servers, set up a local internal NTP server that is used for all internal hosts as a reference timekeeper, and have only the internal NTP server make requests to the public servers. You should ensure that you are standardizing on UTC time across all systems because it will make researching attacks and issues more relevant.

Another option is to consider switching to Network Time Security (NTS). This protocol is a security extension for time protocols, currently focuses on NTP in unicast mode, and operates on TCP port 4460. NTS provides strong cryptographic protection against packet manipulation, prevents tracking, scales, is robust against packet loss, and minimizes the loss of accuracy due to the securing process. Cloudflare is a major provider of NTS servers.

Email and Web

Most email and web protocols have similarities to transport SSL, which is now deprecated by the IETF, replaced by Transport Layer Security (TLS), based on SSL v3.0, symmetric encryption, message integrity (message authentication code), and authentication (PKI digital certificates). Forward secrecy ensures that any future disclosure of encryption keys cannot be used to decrypt any TLS communications previously recorded. HTTPS uses SSL/TLS for secure web-based communications, X.509 digital certificates, and 256-bit encryption keys. Table 17-2 outlines the commonly used email and web protocol port numbers.

Table 17-2 Email/Web Protocol Ports

Port Number	Description	When to Use It
25 TCP	Standard SMTP port	Often blocked to reduce spam
587 TCP	Outgoing SMTP mail port (TLS/Start TLS Port); used by various mail servers for relaying outgoing mail as a modern alternative to port 25. Gmail, Apple MobileMe Mail, and Yahoo SMTP server use this port. See RFC 2476.	Best choice for SMTP (sending messages)
110 TCP	Insecure POP3	Should not use
143 TCP	Insecure IMAP	Should not use
443 TCP, UDP	Hypertext over TLS/SSL	All websites
993 TCP	Secure IMAP (IMAPS)	All IMAP clients
995 TCP	Secure POP3 (POP3S)	All POP3 clients

File Transfer

Secure File Transfer Protocol, or SFTP, is a network protocol that provides file access, file transfer, and file management over any reliable data stream. It was designed by the Internet Engineering Task Force as an extension of the SSH version 2.0 protocol to provide secure file transfer capabilities. The IETF Internet Draft states that, even though this protocol is described in the context of the SSH-2 protocol, it can be used in a number of different applications, such as secure file transfer over TLS and transfer of management information in VPN applications. This protocol assumes that it is run over a secure channel, such as SSH; that the server has already authenticated the client; and that the identity of the client user is available to the protocol. Secure file transfer can be done via SSH, SFTP, and Secure Copy Protocol (SCP) using TCP port 22.

Directory Services

Lightweight Directory Access Protocol (LDAP) is a standard communications protocol used to read data from and write it to Active Directory. Some applications use LDAP to add, remove, or search users and groups in Active Directory or to transport credentials for authenticating users in Active Directory. Every LDAP communication includes a client (such as an application) and a server (such as Active Directory). LDAP over SSL/TLS is also known as LDAPS. With LDAPS, you can improve security across the wire. You can also meet compliance requirements by encrypting all communications between your LDAP-enabled applications. LDAPS operates on port 636.

Remote Access

Secure remote access refers to any security policy, solution, strategy, or process that exists to prevent unauthorized access to your network, its resources, or any confidential or sensitive data. In remote access, secure remote access is a mix of security strategies and not necessarily one specific technology like a VPN. A VPN provides privacy and security to users by creating a private network connection across a public network connection. Remote workers use VPNs to access company networks protected by firewalls that deny inbound and outbound access to systems on the network in a secure manner. VPN services are the most common and well-known form of secure remote access. Multiple protocols can establish VPN connections, such as PPTP, L2TP, and IPsec. Table 17-3 shows the technology utilized in each of these protocols.

Table 17-3 Remote Access Protocols

	L2TP	SSTP	PPTP	Open VPN	L2TP v3	EtherIP	SoftEther VPN
Upper Protocol	IP	IP	IP	Ethernet	Ethernet	Ethernet	Ethernet
Transport Protocol	IPsec	HTTPS	GRE	Specific TCP/UDP	IPsec	IPsec	HTTPS
Session Key	256-bit	256-bit	128-bit	256-bit	256-bit	128-bit	128-bit

Domain Name Resolution

Domain Name System (DNS) is used for resolving IP addresses to names. This is where you often hear Domain Name Resolution, which is the function of DNS. DNS on its own is considered an insecure protocol, and its successor, DNSSEC, is a more secure option. There are still ways to enhance the security of plain old DNS.

You can start by using caching-only DNS servers, using DNS forwards, and using DNS advertisers and resolvers. You should also protect DNS from cache pollution, enable DDNS for secure connections only, disable zone transfers, and use firewalls to control communication to and from the DNS servers.

DNSSEC allows you to verify the responses that you're getting from a DNS server. You can ensure that the DNS request is coming from the correct origin, and you can make sure the information that you're receiving is exactly what was sent from the DNS server. DNSSEC does this using public key cryptography. The records within DNS are effectively digitally signed by a third party. The signed DNS records are stored on the DNS server, and when you receive the DNS information, you can verify the digital signature. If you're a network administrator and you need to manage a switch or router, you can communicate to those devices over a command line using SSH. DNS and DNSSEC use TCP 53, and a lot of standard DNS traffic uses UDP 53.

Routing and Switching

Routing is a process in which a Layer 3 device chooses the best path between the source and destination network. Dynamic routing protocols are used to decrease administrator overhead by reducing the amount of required configuration; however, by default, all the unencrypted routing information is visible to all interested parties, which makes it vulnerable to an attack. You can secure the routing protocols like EIGRP and OSPF by enabling the feature, putting an authentication key chain in place, and applying it to the interface on which you are advertising the routes. You can manage switches and routers via the SNMPv3 previously discussed in this chapter. SNMP utilizes UDP ports 161 and 162.

Network Address Allocation

Network address allocation is done to allocate (assign) IP addresses on networks where there are lots of hosts. We call this Dynamic Host Configuration Protocol (DHCP), which autoassigns IP addresses and uses network subnets to segregate multiple hosts and control network traffic. DHCP uses UDP port 67 as the destination port of a server and UDP port 68 by the client. DHCP operates with options. For example, option 1 is the subnet mask to be applied to the IP address when it's being handed out; if one is not defined, it defaults to the class of the IP address. Option 3 is the default router, or gateway of last resort. Option 6 is which DNS servers to include when assigning the IP address, and finally Option 51 is the lease time for the IP address—that is, how long the DHCP server and client agree to reserve the IP address for a particular client. There are many other options including ones for time, voice and video systems, wireless access points and controllers, and other devices.

Table 17-4 outlines the IPv4 address classes.

Table 17-4 IPv4 Address Classes

Class	Theoretical Address Range	Binary Start	Used for
A	0.0.0.0 to 127.255.255.255	0	Very large networks
B	128.0.0.0 to 191.255.255.255	10	Medium networks
C	192.0.0.0 to 223.255.255.255	110	Small networks
D	224.0.0.0 to 239.255.255.255	1110	Multicast
E	240.0.0.0 to 255.255.255.255	1111	Experimental (reserved)

Subscription Services

Subscription services can be software as a service (SaaS), such as cloud email, Google Gmail, or Microsoft Office 365, as well as network defenses like firewalls, web application filtering, and patching. All these types of subscription services provide patching and updates as part of the purchase of their services, where you technically never own the app.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 17-5 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 17-5 Key Topics for Chapter 17

Key Topic Element	Description	Page Number
Section	Domain Name System Security Extensions (DNSSEC)	426
Section	SSH	427
Section	Secure/Multipurpose Internet Mail Extensions (S/MIME)	428
Section	Secure Real-Time Transport Protocol (SRTP)	430
Section	Lightweight Directory Access Protocol over SSL (LDAPS)	432

Key Topic Element	Description	Page Number
Section	File Transfer Protocol, Secure (FTPS)	432
Section	SSH File Transfer Protocol (SFTP)	434
Section	Simple Network Management Protocol Version 3 (SNMPv3)	434
Section	Hypertext Transfer Protocol over SSL/TLS (HTTPS)	436
Section	IPsec	437
Section	Post Office Protocol (POP)/Internet Message Access Protocol (IMAP)	438
Section	Use cases to implement security for voice and video, time synchronization, and other services	439

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

DNS Security Extensions (DNSSEC), Secure Shell (SSH), Secure/Multipurpose Internet Mail Extensions (S/MIME), Secure Real-Time Transport Protocol (SRTP), Lightweight Directory Access Protocol Over SSL (LDAPS), File Transfer Protocol, Secure (FTPS), Secure File Transfer Protocol (SFTP), Simple Network Management Protocol version 3 (SNMPv3), Hypertext Transfer Protocol Secure (HTTPS), IPsec, Authentication Header (AH), Encapsulating Security Payload (ESP), Tunnel mode, Transport mode, Internet Message Access Protocol (IMAP), Post Office Protocol (POP)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is a secure protocol?
2. How does SSH help secure connections?
3. What cryptography method does S/MIME use?
4. Secure Real-Time Transport Protocol uses which cipher by default?
5. LDAPS is a secure version of LDAP that is used to communicate with Active Directory. What TCP port does LDAPS over SSL/TLS use?



This chapter covers the following topics related to Objective 3.2 (Given a scenario, implement host or application security solutions) of the CompTIA Security+ SY0-601 certification exam:

- Endpoint Protection
  - Antivirus
  - Anti-malware
  - Endpoint detection and response (EDR)
  - DLP
  - Next-generation firewall (NGFW)
  - Host-based intrusion prevention system (HIPS)
  - Host-based intrusion detection system (HIDS)
  - Host-based firewall
- Boot Integrity
  - Boot security/Unified Extensible Firmware Interface (UEFI)
  - Measured boot
  - Boot attestation
- Database
  - Tokenization
  - Salting
  - Hashing
- Application Security
  - Input validations
  - Secure cookies
  - Hypertext Transfer Protocol (HTTP) headers
  - Code signing
  - Allow list
  - Block list/deny list
  - Secure coding practices
  - Static code analysis
    - Manual code review
    - Dynamic code analysis
    - Fuzzing
- Hardening
  - Open ports and services
  - Registry
  - Disk encryption
  - OS
  - Patch management
    - Third-party updates
    - Auto-update
- Self-encrypting drive (SED)/full-disk encryption (FDE)
  - Opal
- Hardware root of trust
- Trusted Platform Module (TPM)
- Sandboxing

Implementing Host or Application Security Solutions

When applications are not correctly programmed, bad things happen; for example, applications can be exploited, allowing attackers in through your creation. No amount of network hardening, user training, or auditing is going to help. Application security is critical for the long-term operations and survival of any business. Application security begins with a secure design and secure coding, from start to finish. This process needs to be maintained over the lifecycle of the application through consistent patching and testing.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 18-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 18-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Endpoint Protection	1–4
Boot Integrity	5
Database	6
Application Security	7
Hardening	8
Self-Encrypting Drive/Full-Disk Encryption	9
Hardware root of trust	10
Trusted Platform Module	11
Sandboxing	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Just how destructive are viruses?
    a. Viruses have cost companies billions.
    b. Viruses have cost companies hundreds of dollars.
    c. Viruses have cost companies millions of dollars.
    d. Viruses are mostly an annoyance and cost very little.

2. What is one of the ways that antimalware software detects malware and hostile code?
    a. Scanning ports of the malware to determine if it is hostile
    b. Making behavior-based observations
    c. Copying the suspect file and compressing it
    d. Realizing that malware is mostly an annoyance

3. What is one of the primary functions of endpoint detection and response (EDR)?
    a. Securing the endpoint by disabling the EDR function
    b. Scanning the ports of the system to determine if the open ports are listening
    c. Providing forensics and analysis tools to research identified threats.
    d. Providing endpoint integrity by encrypting the hard disk

4. What is one of the ways that DLP protects from data loss and misuse?
    a. Monitoring and controlling endpoint activities
    b. Locking data in a bitwise vault
    c. Observing behavior-based activities
    d. Stopping malware from executing by placing it in a sandbox

5. What new specification does Unified Extensible Firmware Interface (UEFI) bring that wasn't in the standard BIOS?
   a. Lowers the threshold for encrypting hard disks
   b. Increases hard disk partition size
   c. Provides extra fields to annotate the type of drive
   d. Stops viruses from executing during shutdown

6. What is the tokenization process?
   a. Using a single token to encrypt all hard disks
   b. Turning nonsensitive data into encrypted data
   c. Turning insensitive data into sensitive data
   d. Turning sensitive data into nonsensitive data

7. What do secure cookies store?
   a. Information about all the sites a user visited over 24 hours
   b. All of the cookies that were collected over two days
   c. Information about the application session after users log out
   d. Information about a user session after the user logs in to an application

8. What process does patch management help users accomplish?
   a. Helps acquire, test, and install multiple patches (code changes) on existing applications
   b. Manages the deployment of security codes to systems
   c. Helps remove patches that do not match the operating system
   d. Manages the Key repository store in the patch management system

9. A self-encrypting drive (SED) installed into a mixed-disk configuration or a configuration containing unencrypted drives operates in which manner?
   a. Operates as an SE disk
   b. Operates as a encrypted disk
   c. Operates as an unencrypted disk
   d. Rejects the drive and powers down

10. As part of the root of trust, any code from outside a system that is intended to run on a secure CPU requires which component?
    a. A signed certificate from a CA that is applied to the CPU
    b. Dedicated RAM that can be accessed by any root authority
    c. Dedicated ROM that can be accessed by the hardware root of trust
    d. Validated code that is secure and signed by a CSR/root

11. UEFI is a replacement for a standard BIOS, UEFI, and the Trusted Platform Module. TPM uses what to sign the log recorded by the UEFI?
    a. Nothing; it's automatically signed each time it is booted.
    b. It uses an enforced signature code and entry.
    c. It uses four elements that include the log, the UEFI, the system time, and a binary signature.
    d. It uses a unique key to digitally sign the log.

12. Sandboxing is a strategy that isolates a test environment for applications to protect them from what?
    a. Malware and viruses
    b. Application information spillage and unauthorized disclosure
    c. Enforcement of access to system and startup log files
    d. Sandthrow attack elements that facilitate hacker access

Foundation Topics

Endpoint Protection

Endpoint protection is a term often used interchangeably with *endpoint security*. Endpoint protection describes security solutions that address endpoint device security issues, securing and protecting endpoints against zero-day exploits, attacks, and inadvertent data leakage resulting from human error.

Targeted attacks and advanced persistent threats (APTs) can't be prevented through antivirus solutions alone, making endpoint protection a necessary component of the full spectrum of security solutions capable of securing data for the world's leading enterprises. Endpoint protection solutions provide centrally managed security solutions that protect endpoints such as servers, workstations, and mobile devices that are used to connect to enterprise networks.

NOTE The idea of data-centric intelligence is to protect the data on the go. For instance, data can be accessed from remote locations; however, the storage of data is not allowed. You must keep a check on the data that comes in and goes out of the organization.

TIP Security experts must place greater emphasis on protecting data. Technology now allows it to go everywhere.

Key Topic

Antivirus

Antivirus software, or *anti-virus software*, also known as antimalware, is a computer program used to prevent, detect, and remove malware. Antivirus software was originally developed to detect and remove computer viruses, hence the name. Just how destructive are viruses? According to a report published by Cybersecurity Ventures in May 2019, the damages from ransomware cost businesses an astonishing $11 billion in lost productivity and remediation. There are many examples of viruses and Trojans damaging networks and costing billions.

Antivirus software is now an integral part of most corporate computing environments and policies. The process of selecting antivirus software varies greatly from organization to organization, but it's important to implement and monitor the solution after it's in place. Every endpoint device, workstation, tablet, and phone should have antivirus software installed. Newer firewalls, routers, and switches implement and integrate antivirus software and should be enabled where possible. Your entire server environment should have antivirus protection installed.

Key Topic Antimalware

In the past, antivirus software typically dealt with older, more well-known threats, such as Trojans, viruses, keyloggers, and worms. Antimalware, on the other hand, emerged to focus on newer, increasingly dangerous threats and infections spread via malvertising and zero-day exploits. Today, however, antivirus and antimalware products are generally the same. Some security vendors continue to refer to their products as *antivirus software* even though their technology is more similar to *antimalware* and covers a wide variety of newer threats.

Antimalware software uses three strategies to protect systems from malicious software, starting with signature-based malware detection, behavior-based malware detection, and sandboxing. These techniques protect against threats from malware in different ways.

Many antivirus and antimalware tools depend on signature-based malware detection. Malicious software is generally identified by comparing a hash of the suspicious code with a database of known malware hashes. Signature-based detection uses a database of known malware definitions to scan for malware.

When antimalware software detects a file that matches the malware signature, it flags that file as potential malware. The limitation of malware detection based on signatures is that it can only identify known malware.

Antimalware software that uses behavior-based malware detection can detect previously unknown malware and threats by identifying malware based on characteristics and behaviors. This type of malware detection evaluates an object based on its intended actions before it can execute that specific behavior. An object is considered malicious if it attempts to perform an abnormal or unauthorized action. Behavior-based detection in newer antimalware products is sometimes powered by machine learning (ML) algorithms.

Sandboxing offers another way for antimalware software to detect malware (see the "Sandboxing" section later in the chapter). A sandbox is an isolated computing environment developed to run unknown applications and prevent them from affecting the underlying system. Antimalware programs that use sandboxing run suspicious or previously unknown programs in a sandbox and monitor the results. If the malware demonstrates malicious behavior (after it is detonated) that behaves suspiciously, the antimalware terminates it.

Endpoint Detection and Response

Endpoint detection and response, also known as endpoint threat detection and response, is an integrated endpoint security solution that combines real-time continuous monitoring and collection of endpoint data with rules-based automated

response and analysis capabilities. EDR is often used to describe emerging security systems that detect and investigate suspicious activities on hosts and endpoints, employing a high degree of automation (machine learning) to enable security teams to quickly identify and respond to threats.

The primary functions of an EDR system are to

Key Topic

1. Monitor and collect activity data from endpoints that could indicate a threat.
2. Analyze this data to identify threat patterns.
3. Automatically respond to identified threats to remove or contain them and notify security personnel.
4. Provide forensics and analysis tools to research identified threats and search for suspicious activities (similar to threat hunting).

Key Topic

Data Loss Prevention

Data loss prevention, or ***DLP***, is a data protection strategy. It is a set of tools that detect potential data breaches and data exfiltration transmissions. It has capabilities to prevent them by monitoring, detecting, and blocking sensitive data while in use/processing, in motion/transit, and at rest.

DLP software classifies regulated, confidential, and business-critical data and identifies violations of policies defined by organizations or within a predefined policy pack, typically driven by regulatory compliance such as the Health Insurance Portability and Accountability Act of 1996 (HIPAA), Payment Card Industry Data Security Standard (PCI DSS), or General Data Protection Regulation (GDPR). When those violations are identified, DLP enforces remediation with alerts, encryption, and other protective actions to prevent end users from accidentally or maliciously sharing data that could put the organization at risk.

Data loss prevention software and tools monitor and control endpoint activities, filter data streams on corporate networks, and monitor data in the cloud to protect data at rest, in motion/transit, and in use/processing. DLP also provides reporting to meet compliance and auditing requirements and identify areas of weakness and anomalies for forensics and incident response.

Key Topic

Next-Generation Firewall

The term ***next-generation firewall (NGFW)*** refers to the third generation of firewall technology, combining the traditional firewall and other functionality such as filtering, application awareness, deep packet inspection, and intrusion prevention systems. Next-generation firewalls filter network traffic to protect an organization

from internal and external threats and also maintain features of stateful sessions. NGFWs provide organizations with Secure Sockets Layer (SSL)/Transport Layer Security (TLS) inspection, application control, intrusion prevention, and advanced visibility across the entire attack surface. As the threat landscape expands due to co-location and multicloud adoption, and businesses grow to satisfy complex customer needs, traditional firewalls fall further behind, unable to offer protection at scale, and leading to poor user experience and weak security posture. NGFWs not only block malware but also include paths for future updates, giving them the flexibility to evolve with the threat landscape and keep the network secure as new threats arise. Figure 18-1 shows the process of an NGFW:

1. Data is sent to the firewall.
2. The packet is inspected and sent to the web server.
3. The return traffic is inspected and sent back to the requester.

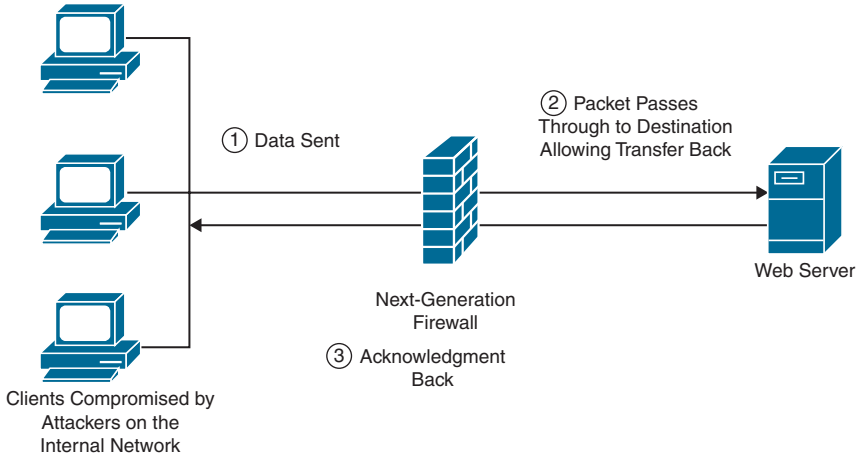


FIGURE 18-1 NGFW Mechanics

Key Topic

Host-based Intrusion Prevention System

The ***host-based intrusion prevention system (HIPS)*** is an intrusion detection system that is capable of monitoring and analyzing the internals of a computing system server as well as the network packets on its network.

Starting from the network layer all the way up to the application layer, HIPS protects from known and unknown malicious attacks. HIPS regularly checks the characteristics of a single host and the various events that occur within the host

for suspicious activities. HIPS can be implemented on various types of machines, including servers, workstations, and computers.

A HIP's typically uses a database of system objects monitored to identify intrusions by analyzing system calls, application logs, and file-system modifications of binaries, password files, capability databases, and access control lists (ACLs). For every object in question, the HIPS remembers each object's attributes and creates a checksum for the contents. This information gets stored in a secure database for later comparison.

The system also checks whether appropriate regions of memory have not been modified. Generally, it keeps a list of trusted programs. A program that oversteps its permissions is blocked from carrying out unapproved actions.

A HIPS has numerous advantages. First, enterprise users require increased protection from unknown malicious attacks. HIPS uses a specific prevention method and system that has a better chance of stopping such attacks as compared to traditional protective measures. A major benefit of using such systems is the requirement to run and manage multiple security applications to protect PCs, such as antivirus, antispyware, and firewalls. Figure 18-2 illustrates the placement of a HIPS in a basic network topology, where the sensor software is typically loaded on servers and PCs to help provide protection and alerting. Every host on a network can become a HIPS sensor; however, sound strategy is to deploy them where critical assets need to be monitored and alerted.

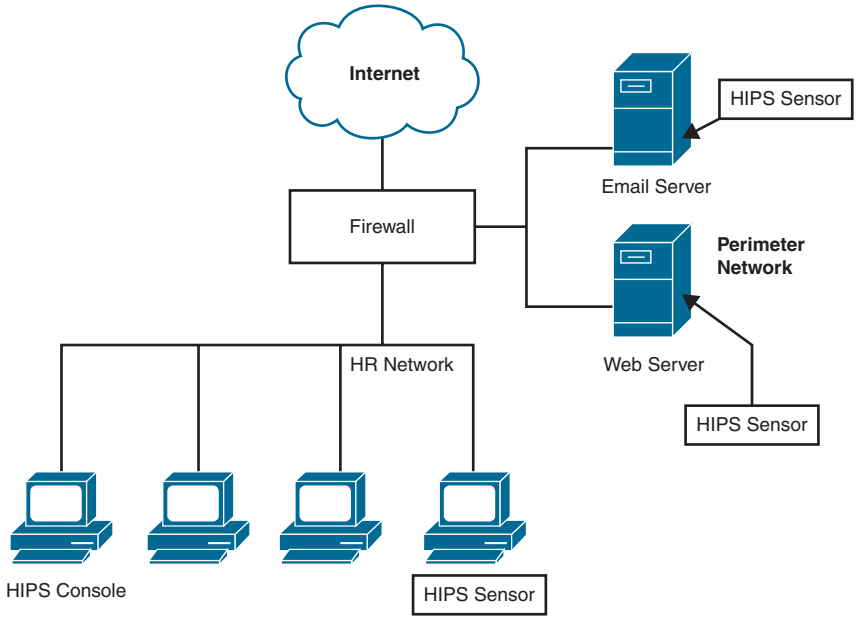


FIGURE 18-2 HIPS and Its Sensor Function in Action

Key Topic Host-based Intrusion Detection System

A ***host-based intrusion detection system (HIDS)*** is an application that operates on information collected from individual computer systems. This vantage point allows a HIDS to analyze activities on the host it monitors at a high level of detail; it can often determine which processes and/or users are involved in malicious activities. HIDSs contextually understand the outcome of an attempted attack because they can directly access and monitor the data files and system processes targeted by these attacks. HIDSs utilize two types of information sources: operating system audit trails and system logs. Operating system audit trails are usually generated at the innermost (kernel) level of the operating system, and typically are more detailed and better protected than system logs. System logs are much less obtuse and much smaller than audit trails, and are normally far easier to comprehend.

Most HIDS software establishes a "digital inventory" of files and their attributes in a known state, and it uses that inventory as a baseline for monitoring any system changes. The "inventory" is usually a file containing all of the SHA-2 checksums for individual files and directories. This must be stored offline on a secured, read-only medium that is not available to an attacker. On a server with no read-only media (a blade server, for example), one method to accomplish this is to store the statically compiled intrusion detection application and its data files on a remote computer. When you wish to run a HIDS report, you can secure copy (SCP) the remote files to /tmp (or its equivalent) on the target server and run them from there. When you modify any files on the server, you should rerun the application and then make a new data set, which should be stored on the remote computer.

The benefit of running HIDS is that it is like a tripwire: if someone enters the system and modifies or changes a file, this activity "trips" HIDS, which alerts the security team. While HIPS can prevent an intrusion, HIDS is only capable of alerting. HIPSs are necessary in any enterprise environment and protect hosts and prevent against known and unknown malicious attacks from the network layer up through the application layer. HIDSs can detect malicious activity and send alerting messages, but they do not prevent attacks. One problem with host-based intrusion detection systems is that any information that they might gather needs to be communicated outside of the machine if a central monitoring system is to be used. If the machine is being actively attacked, particularly in the case of a denial-of-service (DoS) attack, this may not be possible.

Figure 18-3 illustrates how a HIDS works.

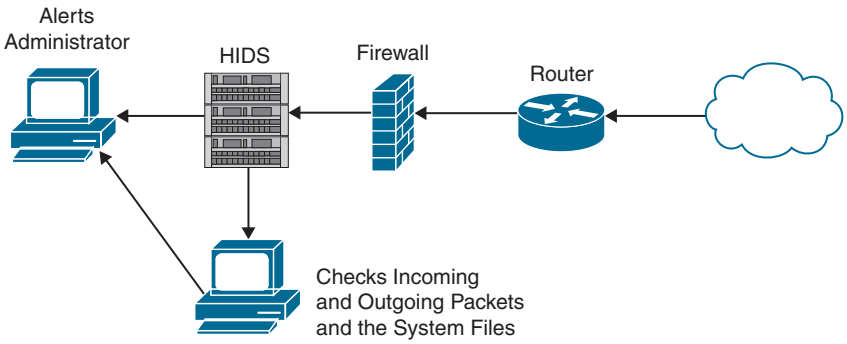


FIGURE 18-3 HIDS in Action

Key Topic

Host-based Firewall

A ***host-based firewall*** is a firewall installed on each individual desktop, laptop computer, or server that controls incoming and outgoing network traffic and determines whether to allow it into a particular device (for example, the Microsoft Windows Defender Firewall that comes with a Windows-based computer). These types of firewalls provide a granular way to protect the individual hosts from viruses and malware and to control the spread of harmful infections throughout the network.

The vast majority of companies use host-based firewalls in addition to perimeter-based firewalls to enhance internal security. For example, some of the malware attacks that might get past a perimeter firewall can be stopped at an individual device or workstation, using a host-based firewall. A host-based firewall setup can also be simpler for some users to implement and manage. The host-based firewall can also be configured to a particular computer, where customization can make the firewall more effective.

Windows Defender Firewall with Advanced Security provides host-based, two-way network traffic filtering and blocks unauthorized network traffic flowing into or out of the local device. You should configure your Windows Defender Firewall utilizing the Microsoft best practice guides specific to your Windows version. To open Windows Defender Firewall with Advanced Security, go to the Start menu, select Run, type **wf.msc**, and then select OK. As you can see from Figure 18-4, Windows Defender Firewall has quite a few configuration options, providing flexibility and functionality.

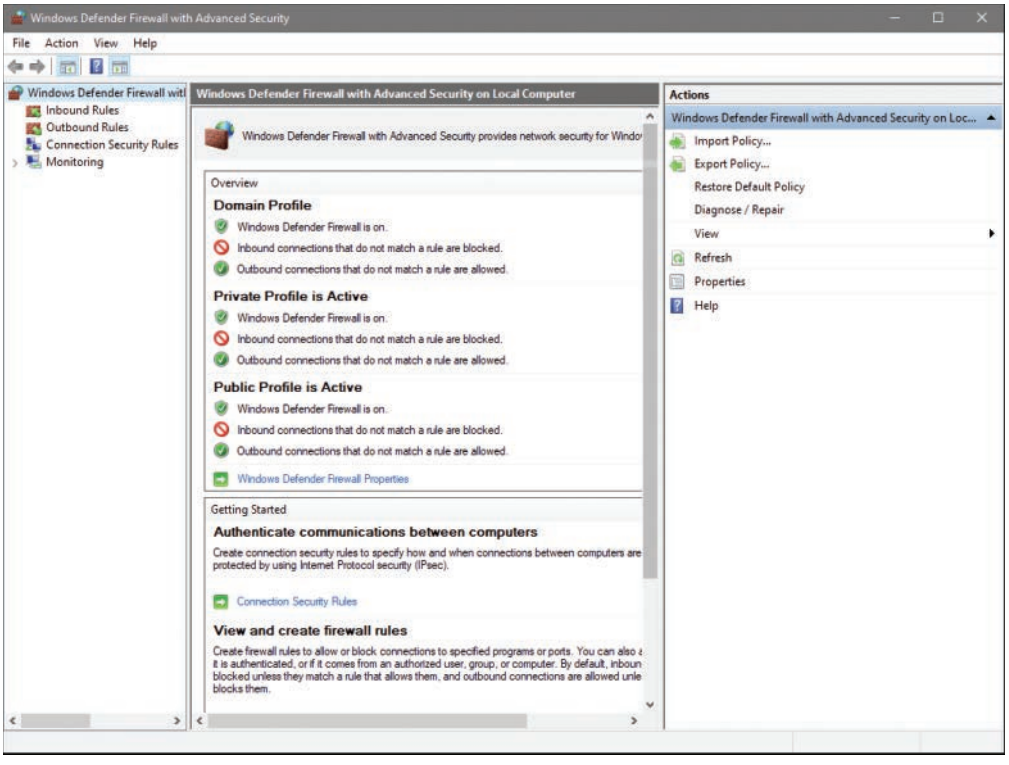


FIGURE 18-4 Windows Defender Firewall

Windows Defender Firewall can be configured to allow specific applications access to the Internet via inbound and outbound rules:

- Rules can be inbound, outbound, or both.
- Rules can include services, which specify the type of traffic or port number.
- Rules can be set to allow or deny.
- Most firewall configurations start by blocking all traffic by default and then allowing only specific traffic in and out.

Key Topic

Boot Integrity

Boot integrity refers to using a secure method to boot a system and verify the integrity of the operating system and loading mechanism. There is actually a boot integrity usage model created by MITRE, and it goes beyond just computers. Boot integrity represents the first step toward achieving a trusted infrastructure. This

model applies equally well to the compute, network, and storage domains. Every network switch, router, or firewall runs a compute layer operating a specialized operating system to provide networking and security functions. This model enables a service provider to make claims about the boot integrity of the network and compute platforms, as well as the operating system and hypervisor instances running in them. Boot integrity supported in the hardware makes the system robust and less vulnerable to tampering and targeted attacks. It enables an infrastructure service provider to make quantifiable claims about the boot-time integrity of the prelaunch and launch components. This provides a means, therefore, to observe and measure the integrity of the infrastructure. In a cloud infrastructure, these security features refer to the virtualization technology in use, which comprises two layers:

- The boot integrity of the BIOS, firmware, and hypervisor. This capability can be referred to as a trusted platform boot.
- The boot integrity of the virtual machines that host the workloads and applications. You want these applications to run on trusted virtual machines.

Key Topic

Boot Security/Unified Extensible Firmware Interface

Unified Extensible Firmware Interface (UEFI) is a specification for a software program that connects a computer's firmware to its operating system. UEFI is expected to eventually replace the BIOS. If you purchase a new computer today, it likely has UEFI.

Like the BIOS, UEFI is installed at the time of manufacturing and is the first program that runs when a computer is turned on. It checks to see what hardware components the computing device has, wakes up the components, and then hands them over to the operating system. This newer specification addresses several limitations of the BIOS, including restrictions on hard disk partition size and the amount of time the BIOS takes to perform its tasks.

Because UEFI is programmable, original equipment manufacturer (OEM) developers can add applications and drivers, allowing UEFI to function as a lightweight operating system.

Unified Extensible Firmware Interface is managed by a group of chipset, hardware, system, firmware, and operating system vendors called the UEFI Forum. The specification is most often pronounced by naming the letters *U-E-F-I*.

Measured Boot

Measured boot is a feature that was introduced in Windows 8; it was created to help better protect your machine from rootkits and other malware. Measured boot

checks each startup component, including the firmware all the way to the boot drivers, and it stores this information in what is called a Trusted Platform Module (TPM). The PC's firmware logs the boot process, and Windows can send it to a trusted server that can objectively assess the PC's health.

In Windows 10, measured boot uses the following process:

Key Topic

1. The PC's UEFI firmware stores a hash of the firmware, bootloader, boot drivers, and everything that will be loaded before the antimalware app in the TPM.
2. At the end of the startup process, Windows starts the non-Microsoft remote attestation client.
3. The TPM uses the unique key to digitally sign the log recorded by the UEFI.
4. The client sends the log to the server, possibly with other security information.

Figure 18-5 illustrates the measured boot and remote attestation process.

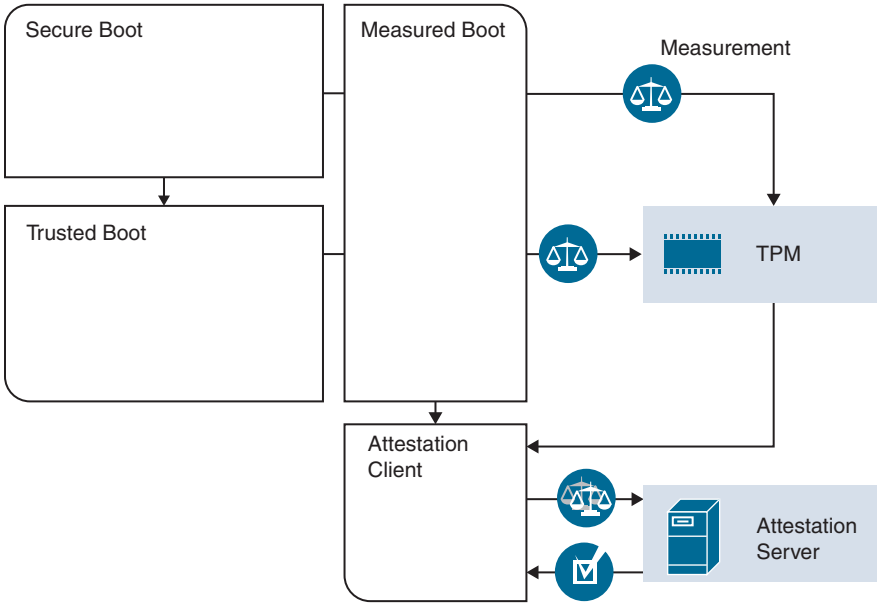


FIGURE 18-5 Measured Boot

Key Topic

Boot Attestation

Boot attestation enables a remote platform to measure and report its system state in a secure way to a third party. In boot attestation, software integrity measurements

are immediately committed to during boot, thus relaxing the traditional requirement for secure storage and reporting as shown previously in Figure 18-5. The third-party platform attests to a verifying platform about the trustworthiness of the software running on the host platform done in the post-boot process.

Key Topic Database

Securing databases and applications that utilize these databases starts with five best practices:

- Separate the database and web servers.
- Encrypt stored data, files, and backups of the database.
- Use a web application (or database) firewall (WAF). Remember to keep patches current and enable security controls.
- Ensure physical security (this is usually a given, but don't leave anything to chance). Ensure your databases are in locked cabinets, utilize the hardening best practices for your specific database, manage access to the database, and tightly guard secrets.
- Make sure you have audit procedures and active monitoring of the database activity enabled. Also, ensure alerts generate tickets and receive proper attention.

Tokenization

Database ***tokenization*** is the process of turning sensitive data into nonsensitive data called *tokens* that can be used in a database or internal system without bringing it into scope. The tokens are then sent to an organization's internal systems for use, and the original data is stored in a secure token vault. There is no key, or algorithm, that can be used to derive the original data for a token; instead, tokens are sent to a token vault. Figure 18-6 illustrates the tokenization process.

The token value can be used in applications as a substitute for the real data. If and when the real data needs to be retrieved, the token is submitted to the vault, and the index is used to cross reference and fetch the real value for use in the authorization process. To the end user, this operation is performed seamlessly by the browser or application nearly instantaneously. Users are likely not even aware that the data is stored in the cloud in a different format.

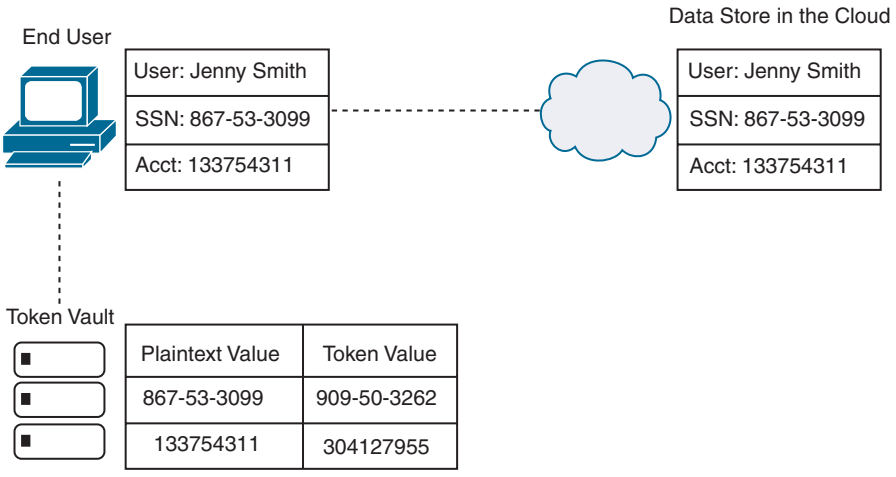


FIGURE 18-6 Tokenization

The advantage of tokens is that there is no mathematical relationship to the real data they represent. If they are breached, they have no meaning. No key can reverse them back to the real data values. Consideration can also be given to the design of a token to make it more useful. Apple Pay uses a similar technology; for example, the last four digits of a payment card number can be preserved in the token so that the tokenized number (or a portion of it) can be printed on the customer's receipt so he or she can see a reference to his or her actual credit card number. The printed characters might be all asterisks plus those last four digits. In this case, the merchant has only a token, not a real card number, for security purposes.

NOTE Remember that tokenization assigns a random surrogate value with no mathematical relationship and can still be reversed by linking the token back to the original data. Outside of the system, a token has no value; it is just meaningless data.

Salting

In cryptography, a ***salt*** is random data that is used as an additional input to a one-way function that hashes data, a password, or a passphrase. Salts are used to safeguard passwords in storage. Historically, a password was stored in plaintext on a system, but over time additional safeguards were developed to protect a user's password against being read from the system. A salt is one of those methods.

Salts can help defend against a precomputed hash attack (for example, rainbow tables). Because salts do not have to be memorized by humans, they can make the size of the hash table required for a successful attack prohibitively large without placing a burden on the users. Because salts are different in each case, they also protect commonly used passwords, or those users who use the same password on several sites, by making all salted hash instances for the same password different from each other.

Key Topic

Hashing

Hashing is a one-way function where data is mapped to a fixed-length value. Hashing is primarily used for authentication. Salting is an additional step that can be added during hashing (typically seen in association with hashed passwords); it adds an additional value to the end of the password to change the hash value produced. Whereas encryption is meant to protect data in transit, hashing is meant to verify that a file or piece of data hasn't been altered—that it is authentic. In other words, it serves as a checksum.

Here's how hashing works. Each hashing algorithm outputs at a fixed length. For instance, let's look at SHA-256, which means that the algorithm is going to output a hash value that is 256 bits, usually represented by a 64-character hexadecimal string. The size of the data block differs from one algorithm to another. But for a particular algorithm, it remains the same. For example, SHA-1 takes in the message/data in blocks of 512 bits only. So, if the message is exactly 512 bits long, the hash function runs only once, which is 80 rounds in case of SHA-1. Similarly, if the message is 1024 bits, it's divided into two blocks of 512 bits, and the hash function is run twice. Because there is nearly no chance of a message being exactly the same size as the block, a technique called *padding* is used. Padding takes the entire message and divides it into fixed-size data blocks. The hash function is repeated as many times as the number of data blocks. The output of the first data block is fed as input along with the second data block. If you change one bit anywhere in the message, the entire hash value changes; this is called the *avalanche effect*.

Every hash value is unique. If two different files produce the same unique hash value, this is called a collision, and it makes the algorithm essentially useless. Let's say you want to digitally sign a piece of software and make it available for download on your website. To do this, you create a hash of the script or executable you're signing. Then after adding your digital signature, you hash that too. Following this, the whole thing is encrypted so that it can be downloaded.

Key Topic

Application Security

Application security describes security measures at the application level, the goal of which is to prevent data or code within the app from being stolen, intercepted, or

hijacked. It encompasses the security considerations that happen during application development and design, and it also involves systems and approaches to protect apps after they get deployed.

Application security can include hardware, software, and procedures that identify or minimize security vulnerabilities. A router that prevents anyone from viewing a computer's IP address from the Internet is a form of hardware application security. But security measures at the application level are also typically built into the software, such as an application firewall, that strictly defines what activities are allowed and prohibited. Procedures can entail things like an application security routine that includes protocols such as regular testing.

Key Topic

Input Validations

Input validation is the first step in checking the type and content of data supplied by a user or application. Improper input validation is a major factor in many web security vulnerabilities, including cross-site scripting (XSS) and SQL injection.

What is input validation? Any system or application that processes input data needs to ensure that it is valid. This applies both to information provided directly by the user and data received from other systems. Validation can be done on many levels, from simply checking the input types and lengths (syntactic validation) to ensuring that supplied values are valid in the application context (semantic validation).

In web applications, input validation typically means checking the values of web form input fields to ensure that a date field contains a valid date, an email field contains a valid email address, and so on. This initial client-side validation is performed directly in the browser, but submitted values also need to be checked on the server side.

NOTE While we generally talk about *user input* or *user-controlled input*, a good practice is to check all inputs to an application and treat them as untrusted until validated.

Let's see how to ensure proper input validation in web applications. Traditionally, form fields and other inputs were validated in JavaScript, either manually or using a dedicated library. Implementing validation is a tedious and error-prone process, so it's a good idea to check for existing validation features before you go the DIY route. Many languages and frameworks come with built-in validators that make form validation much easier and more reliable. For input data that should match a specific JSON or XML schema, you should validate input against that schema.

Secure Cookies

The ***secure cookie*** attribute is an option that can be set by the application server when sending a new cookie to the user within an HTTP response. The purpose of the secure attribute is to prevent cookies from being observed by unauthorized parties due to the transmission of the cookies in clear text.

Cookies may contain sensitive information that shouldn't be accessible to an attacker eavesdropping on a channel. To ensure that cookies aren't transmitted in clear text, you can send them with a secure flag.

Web browsers supporting the "secure" flag only send cookies having the secure flag when the request uses HTTPS. This means that setting the secure flag of a cookie prevents browsers from sending it over an unencrypted channel. The unsecure cookies issue is commonly raised in penetration test reports performed by pen testers if the environment they are running on is missing the correct credentials.

The secure session cookies store information about a user session after the user logs in to an application. This information can be highly sensitive because an attacker can use a session cookie to impersonate the victim in what's called *session hijacking* or *cookie hijacking*. The goal of session hijacking is to steal a valid and authorized cookie from real users to gain unauthorized access to their account or an entire system.

Hypertext Transfer Protocol Headers

Hypertext Transfer Protocol (HTTP) is an application-layer protocol for transmitting hypermedia documents, such as HTML. It was designed for communication between web browsers and web servers, but it can also be used for other purposes. HTTP (port 80) has mostly been replaced with HTTPS (port 443). As a matter of fact, Chrome and Firefox can alert you to an untrusted site (displaying an "untrusted" white page if you browse to an HTTP-only page).

HTTP headers can be grouped according to their contexts, and there are hundreds of header and header fields. Here, we touch on only the most common and relevant:

- General headers apply to both requests and responses, but with no relation to the data transmitted in the body.
- Request headers contain more information about the resource to be fetched or about the client requesting the resource.
- Response headers hold additional information about the response, like its location or about the server providing it.
- Entity headers contain information about the body of the resource, like its content length or MIME type.

End-to-End Headers

End-to-end headers must be transmitted to the final recipient of the message: the server for a request or the client for a response. Intermediate proxies must retransmit these headers unmodified.

Hop-by-Hop Headers

Hop-by-hop headers are meaningful only for a single transport-level connection and must not be retransmitted by proxies or cached. Note that only hop-by-hop headers may be set using the Connection general header.

One of the more common headers is User-Agent, which contains a characteristic string that enables the network protocol peers to identify the application type, operating system, software vendor, or software version of the requesting software user agent—with the web browser showing if you're on a Mac or Linux host or using your cell phone. That's how it is able to render a page to your phone properly. Because HTTP traffic is in plaintext, that is, clear text, you need to protect your server with firewalls or HIDS to alert and mitigate potential threats.

The latest version of HTTP is HTTP/3, the third major version. HTTP/3 uses QUIC, a transport layer protocol, where user space congestion control is used over UDP.



Code Signing

Code signing is the process of digitally signing executables and scripts to confirm the software author can guarantee that the code has not been altered or corrupted since it was signed. The process employs the use of a cryptographic hash to validate authenticity and integrity. Digitally signing code provides both data integrity to prove that the code was not modified and source authentication to identify who signed the code. Furthermore, digitally signing code defines code signing use cases and can help identify security problems that arise when applying code signing solutions to those use cases. Code signing is what allows you to be sure you are downloading the right file from the right author/publisher instead of from an attacker who wants to steal your data.

Before developers can sign their work, they need to generate a public/private key pair. This is often done locally through software tools such as OpenSSL. Developers then give the public key and the organization's identity information to a trustworthy certificate authority (CA). The CA verifies the authenticity of identity information and then issues the certificate to the developer. This is the code signing certificate that was signed by the CA's private key and contains the developer organization's identity and the developer's public key.

When developers are ready to "sign" their work to establish authorship, they take all the code they wrote and then hash it. The value that is spit out is then encoded using a private key (usually generated by the author), along with the code signing certificate that contains the public key and identity of the author (proving the authorship). The output of this process is then added to the software to be shipped out.

This process constitutes a code signing operation. The public key of the CA is already preinstalled in most browsers and operating system trust stores. When a user tries to download the software, that user uses the CA's public key to verify the authenticity of the code signing certificate embedded in the software to confirm that it's from a trustworthy CA. The developer's public key is then extracted from the certificate and used to decrypt the encrypted hash. Then the software is hashed again, and the new value is compared to the decrypted one. If the user's hash value and developer's hash value match, the software hasn't been corrupted or tampered with during transmission. The result is your assurance that you can run the code safely.



Allow List

The ***allow list***, or whitelist, is much easier and safer to use for well-defined inputs such as numbers, dates, or postcodes. That way, you can clearly specify permitted values and reject everything else. With HTML5 form validation, you get predefined allow list logic in the built-in data type definitions, so if you indicate that a field contains an email address, you have ready email validation. If only a handful of values are expected, you can use regular expressions to explicitly allow list them.

The allow list gets tricky with free-form text fields, where you need some way to allow the vast majority of available characters, potentially in many different alphabets. Unicode character categories can be useful to allow, for example, only letters and numbers in a variety of international scripts. You should also apply normalization to ensure that all input uses the same encoding and no invalid characters are present. The allow list requires a lot of resource time to maintain and needs to be continuously updated as the company works with new applications and removes old ones.



Block List/Deny List

When approaching input validation from a security perspective, you might be tempted to implement it by simply disallowing elements that might be used in an injection attack. For example, you might try to ban apostrophes and semicolons to prevent SQL injection, parentheses to stop malicious users from inserting a JavaScript function, or angle brackets to eliminate the risk of someone entering HTML

tags. This is called ***block listing*** or ***deny listing***, and it's usually a bad idea because the developer can't possibly know or anticipate all possible inputs and attack vectors. Blocklist-based validation is hard to implement and maintain and very easy for an attacker to bypass.

Let's say you want to use these lists anyway. Now you have added an additional maintenance point, and it's required that you understand these block lists can potentially break things, and your upper layer programming should not depend on it to stop attacks.

Secure Coding Practices

Secure coding practices must be incorporated into all lifecycle stages of an application development process. The software development lifecycle (SDLC) processes incorporate major components of a development process, starting with requirements, architecture design, implementation, testing, deployment, and maintenance. You can integrate secure coding principles into SDLC components by providing a general description of how the secure coding principles are addressed in architecture and design documents. If a secure coding principle is not applicable to your project, this point should be explicitly documented along with a brief explanation. You can perform automated and static application security testing as part of the overall application testing process to ensure checkpoints are met through the entire process. Development and testing environments should redact all sensitive data or use deidentified data. Figure 18-7 illustrates security processes in the SDLC.

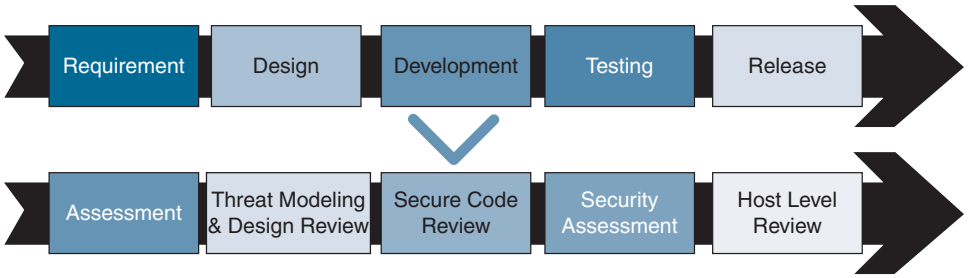


FIGURE 18-7 Security Processes in the SDLC

Static Code Analysis

Static application security testing (SAST), or ***static code analysis***, is a testing methodology that analyzes source code to find security vulnerabilities that make your organization's applications susceptible to attack. SAST scans an application before the code is compiled. It's also known as *known environment/white box testing*. Static

code analysis takes place very early in the software development lifecycle because it does not require a working application and can take place without code being executed. It helps developers identify vulnerabilities in the initial stages of development and quickly resolve issues without breaking builds or passing on vulnerabilities to the final release of the application.

Static code analysis tools give developers real-time feedback while they code and can help them fix issues before they pass the code to the next phase of the SDLC. This analysis prevents costly and security-related issues from being considered an afterthought. SAST tools also provide graphical representations of the issues found, from source to sink. They help you navigate the code easier. Some tools point out the exact location of vulnerabilities and highlight the risky code. Tools can also provide in-depth guidance on how to fix issues and the best place in the code to fix them, without requiring deep security domain expertise.

Developers can also create the customized reports they need with SAST tools; these reports can be exported offline and tracked using dashboards. Tracking all the security issues reported by the tool in an organized way can help developers remediate these issues promptly and release applications with minimal problems. This process contributes to the creation of a secure SDLC.

It's important to note that SAST tools must be run on the application on a regular basis, such as during daily/monthly builds, every time code is checked in, or during a code release.

There are six steps to running a static code analysis effectively:



Step 1. Finalize the tool; select a static analysis tool that can perform the specific code review. The tool should be able to understand the underlying framework of the software.

Step 2. Create the scanning infrastructure and deploy the tool.

Step 3. Customize the tool to suit the needs of the analysis.

Step 4. Prioritize and onboard applications. Scan high-risk applications first.

Step 5. Analyze scan results, remove all false positives, and track each result. The team should be apprised of each defect and schedule appropriate and timely remediation.

Step 6. Provide governance and training. Proper governance ensures that your development team employs the tools properly and consistently.

Manual Code Review

Manual code review is the process of reading source code line by line in an attempt to identify potential vulnerabilities. This is a tedious process that requires skill, experience, persistence, and patience. There are three primary phases of a manual secure code review and it starts with the interview, then the code review, and finally reporting results.

During the interview with the developers, the review team has a chance to understand the intent of the application before reviewing the code. First getting a basic understanding of the intended business function or capability of the application, the interview focuses on key security touch points, such as the developers' approach to authentication, data validation, and logging.

After the interview, the code review begins where the review team works individually to review the application as a whole. Rather than handing off individual code files to specific team members, each member reviews the entire application. This approach plays to the strengths of each individual reviewer, often resulting in the identification of different, yet relevant, findings. The other advantage, of course, is that having multiple eyes on the same set of code serves as a quality check to ensure findings are valid.

After the individual code reviews are completed, the team meets to share results. Each reviewer has the chance to review the others' findings, providing an opportunity to discuss why certain findings may appear in one team member's list but not in another's. This reporting results phase also helps ensure that the findings being reported are relevant. The final list of findings, along with descriptions and potential mitigations, is then presented to the developers using a standard report format.

Dynamic Code Analysis

Dynamic code analysis adopts the opposite approach in regard to static code analysis: it is executed while a program is in operation. Dynamic application security testing (DAST) looks at the application from the outside in, by examining it in its running state and trying to manipulate it to discover security vulnerabilities. The dynamic test simulates attacks against a web application and analyzes the application's reactions, determining whether it is vulnerable. Dynamic analysis is capable of exposing a subtle flaw or vulnerability too complicated for static analysis alone to reveal. Used wisely, automated tools can dramatically improve the return on testing investment. Automated testing tools are an ideal option in certain situations. Automating tests that are run on a regular basis during the SDLC is also helpful. As the enterprise strives to secure the SDLC, it must be noted that there is no panacea. Neither static nor dynamic testing alone can offer blanket protection. Ideally, an

enterprise will perform both static and dynamic analyses. This approach benefits from the synergistic relationship that exists between static and dynamic testing.

Fuzzing

Fuzzing is the art of automatic bug detection involving invalid, unexpected, or random data as inputs to a computer or program. The goal of fuzzing is to stress the application and cause unexpected behavior, resource leaks, or crashes. This capability makes a fuzzer a great asset in assessing the security and stability of applications.

Fuzzing works by first generating test cases. Each security test case can be generated as a random or semirandom data set, and then is sent as input to the application. The data set can be generated either in conformance to the format requirements of the system's input or as a completely malformed chunk of data the system was not meant to understand or process.

What do you think would happen to an application if negative numbers, null characters, or even special characters were sent to some input fields? Do you know how your application would behave? Answering these questions helps with figuring out the trigger. Depending on your needs, you would select the best fuzzer for the job. Mutation or generation fuzzers are defined by the way they handle test case generation. Generation fuzzers generate new test cases from a supplied model, whereas mutation fuzzers mutate a supplied seed input object. There are fuzzers that can do both.

Hardening

Application ***hardening***, also known as application *shielding*, is the act of applying levels of security to protect applications from IP theft, misuse, vulnerability exploitation, tampering, or even repackaging by people with ill intentions. Application hardening is usually performed via security solutions or tools with specialized hardening capabilities that greatly increase the effort required by attackers to modify the application, making it no longer viable or worthwhile to target. The most robust tools shield applications from both static and dynamic threats.

Open Ports and Services

All communication that happens over the Internet is exchanged via ports. Every IP address contains two kinds of ports, TCP and UDP, and there can be up to 65,535 of each for any given IP address. Services that connect to the Internet (web browsers, email clients, and file transfer services, for example) use specific ports to receive information.

Any Internet-connected service requires specific ports to be open in order to function. A problem arises when legitimate services are exploited through code vulnerabilities or malicious services are introduced to a system via malware. Cybercriminals can use these services in conjunction with open ports to gain access to sensitive data.

Closing unused ports is like shutting the door on those cybercriminals. That's why it's considered best practice to close any ports that aren't associated with a known legitimate service.

As system administrator, you can scan for and close open ports that are active and exchanging information on your networks. Closing open ports requires knowing which ports are actually required by the applications and services running on a network. Some of them are universal—for example, port 80 is the port for web traffic (HTTP). Others are reserved by specific services.

Once you know which ports must remain open, you can conduct a scan to identify open ports that might be exposing your systems to cyber attacks. Many free online tools make this scanning process easier. If a port is (1) open and (2) not associated with any known service on the network, you should investigate which application or service is using the port and should remove it immediately. You should also have continuous monitoring in place to ensure an attacker hasn't found a port you missed and has gained a foothold.



Registry

A user ***registry*** is a data source that contains account information, such as usernames and passwords. An application server can retrieve this information to authenticate and authorize users. Securing and hardening of the registry can vary with the operating system, and if possible be done right after the system is built and before it is put into production. Be warned, however, that performing any work on the registry could completely destroy your system, making it irrecoverable. In most cases with the Windows registry, you use an MMC (snap-in) to add the security template. Enter **regedit** from the Start > Run dialog box and you are presented with the registry.

There are several expert guides you can use to harden your specific version of Windows. It is highly recommend you utilize one of them. Best practices for the registry are as follows: do not allow remote access, do *not* store passwords using reversible encryption methods, and restrict the ability to access the computer to only network administrators. You also should deny guest accounts the ability to log on as a service, as a batch job, locally, or via Remote Desktop Protocol (RDP). Then configure the Microsoft network client to always digitally sign communications.

Disk Encryption

Disk encryption is a technology which protects information by converting it into unreadable code that cannot be deciphered easily by unauthorized people. Disk encryption uses disk encryption software or hardware to encrypt every bit of data that goes on a disk or disk volume.

Full-disk encryption is a cryptographic method that applies encryption to the entire hard drive, including data, files, the operating system, and software programs. Full-disk encryption places an exterior guard on the internal contents of the device. Unlike past iterations of full-disk encryption, the process to encrypt hard drives has become quite simple and is supported by all the major vendors.

Under no circumstance should you store confidential files unencrypted (at rest) on your hard disks. You should use disk encryption when storing confidential or sensitive data. Disk encryption can mitigate risks of data exposure from loss or theft of stored data. Full-disk encryption can provide "blanket" protection so users do not have to protect individually stored files, and it ensures that any remnants of data are secure. Such remnants can be temporary files, browser cache files, and application-specific automatic backups. All current operating systems provide a disk encryption capability. All major operating systems support disk encryption, so there should be no reason not to use it.

Operating System

Hardening of the operating system is the act of configuring an OS securely, updating it, creating rules and policies to help govern the system in a secure manner, and removing unnecessary applications and services. This is done to minimize a computer OS's exposure and attack surface to threats and to mitigate possible risk. Although every OS is different, you can do some common things to protect your operating system:

- Remove unnecessary and unused programs. Every program installed on a device is a potential entry point for a bad actor, so be sure to clean them up regularly. If a program has not been okayed or vetted by the company, it should not be allowed. Hackers look for security holes when attempting to compromise networks, so this is your chance to minimize their opportunities.
- Use updates, keeping your programs up to date and installing the latest versions of updates. There's no single action that ensures protection, especially from zero-day attacks, but using updates is an easy and effective step to take.
- Make patch management a part of any regular security regimen. Doing so involves planning, testing, implementing, and auditing consistently to ensure the OS is patched, as well as individual programs on the client's computer.

- Establish and use group policies. Sometimes, user error can lead to a successful cyber attack. One way to prevent such attacks is to define the groups that have access and stick to those rules. Update user policies and make sure all users are aware of and compliant with these procedures. For instance, enforce the use of strong passwords. Also, consider using security templates in the group policy objects (GPOs); these are often used by corporate environments and are essentially text files that represent a security configuration. You could basically use a security template to help manage your group policy and ensure consistency across your entire organization.
- Configure baselines. This is how you measure changes in networking, hardware, software, and so on. Baselines are created by selecting something to measure and doing so consistently for a period of time. After you establish a baseline, measure it on a schedule that meets your security maintenance standards and needs.
- Do not allow users to open web pages from inside applications like Word, Excel, and Outlook. Also, disable macros.

Key Topic

Patch Management

Patch management is a process that helps acquire, test, and install multiple patches (code changes) on existing applications and software tools on a computer, enabling systems to stay updated with existing patches and determining which patches are the appropriate ones. Managing patches thus becomes easy and simple. One of the most simple, common, and overlooked ways to protect systems is to keep them patched. Consider securing your auto-updates with onsite servers. Third-party applications are rarely checked for patches, updates, or whether the application creator discontinued support; therefore, you must stay vigilant.

- **Third-party updates:** Nearly every application, tool, or widget provides updates, and usually on a regular basis. Therefore, you should make sure your users do not disable this capability by using group policies. Set up systems to perform updates during off hours. Keep an eye on updates to make sure they do not cause more harm than good. Generally, it's a good idea to test all third-party patches on a handful of test machines. Provided they perform well, you can roll them out en masse. Some third-party patches do *not* automatically notify or download for users. In that case, you should have a good inventory of every software package running in the organization and monitor them for updates and known exploits.
- **Auto-updates:** Most operating systems enable auto-updates by default, and if you do not force a reboot every evening through group policies, users tend

to hit the "Snooze" button on rebooting for patch updates. You should allow the download of auto-updates of critical patches where possible and schedule the reboot, rollout, or installation during off hours to make sure you have as small of an impact on the organization as possible. First, make sure your backups are working and that they show that they successfully performed a backup recently. You can also configure an onsite patching server and ensure all your systems only connect to your patching server to receive their updates. In addition to saving Internet bandwidth, this approach also allows you to test updates on lab systems to make sure they do not have adverse affects.

Key Topic

Self-Encrypting Drive/Full-Disk Encryption

A ***self-encrypting drive (SED)*** is a disk drive that uses an encryption key to secure the data stored on the disk. This encryption protects the data and array from data theft when a drive is removed from the array.

SED operates across all disks in an array at once. If one drive in a RAID set is removed from the array, a new set of encryption key shares is generated automatically and shared among the remaining disks. If a second drive is removed from the same RAID set, another set of encryption key shares is generated. SEDs are configured at the factory. When the drives are installed into an array, the array automatically detects the new SEDs and locks them. This process is automatic; there are no GUI user controls for SED.

All of the drives in an array, including spares, must be of the same type and model and must be running specific firmware. A self-encrypting drive installed into a mixed-disk configuration or a configuration containing unencrypted drives operates as an unencrypted disk. Likewise, a pool consisting of all SEDs might replicate to a pool with only a few SEDs or no SEDs at all. Hard drive full-disk encryption is made by HDD vendors using the OPAL and Enterprise standards developed by the Trusted Computing Group. Key management takes place within the hard disk controller, and encryption keys are 128- or 256-bit Advanced Encryption Standard (AES) keys. ***Full-disk encryption (FDE)*** and self-encrypting drives encrypt data as it is written to the disk and decrypt data as it is read off the disk. FDE makes sense for laptops, which are highly susceptible to loss or theft, but not so much in a data center.

NOTE Self-encrypting drives are identified in the GUI with a gold key icon.

NOTE The term *self-encrypting drive (SED)* is often used when referring to full-disk encryption (FDE) on hard disks. The Trusted Computing Group (TCG) Opal security subsystem storage standard provides industry-accepted standardized SEDs. SEDs automatically encrypt all data in the drive, preventing attackers from accessing the data through the operating system. SED vendors include Seagate Technology, Hitachi, Western Digital, Samsung, and Toshiba.

OPAL

Opal is a standard developed by the Trusted Computing Group (TCG). Hitachi, Western Digital, Seagate, Samsung, and Toshiba are the disk drive manufacturers offering TCG OPAL SATA drives. The Opal Security Subsystem Class (or SSC) is an implementation profile for storage devices built to protect the confidentiality of stored user data against unauthorized access once it leaves the owner's control (involving a power cycle and subsequent deauthentication). With Opal, you can also enable interoperability between multiple SD vendors within a system.

The Opal SSC encompasses these functions: Security Provider Support Interface Communication Protocol, cryptographic features, authentication, table management, access control and personalization issuance, and SSC discovery. SCC discovery is a process in which the host examines the storage device's configurations, capabilities, and class.

Key Topic Hardware Root of Trust

The ***hardware root of trust*** is the foundation on which all secure operations of a computing system depend. It contains the keys used for cryptographic functions and enables a secure boot process. It is inherently trusted and therefore must be secure by design. Hardware root of trust can be defined by the four basic building blocks:

- The protective hardware provides a trusted execution environment (TEE) for the privileged software to run.
- At a minimum, it must perform one or more proven cryptographic functions, like AES-based cryptography
- A form of tamper protection must be present and available for the entire runtime.
- It must have a flexible yet simple user interface that the host can interact with, through either the host CPU and/or a host controller toggling general-purpose I/O (GPIOs).

To meet the root of trust criteria, a hardware root of trust needs to include a variety of components, starting with a security perimeter. The security perimeter defines what needs to be protected on the system on a chip (SoC). It can be implemented in various ways, including via a private bus that connects to the main bus through a gateway.

Next, a root of trust is required to have a secure CPU that runs secure code software/firmware. The security features supported in a hardware root of trust policy are enabled and used by the software running on that CPU. The resources around the CPU help facilitate the security and performance of these functions.

The third element of a root of trust is the runtime memory. When running software on the CPU, developers and designers need to protect the runtime data required by the software—more explicitly, the STACK, HEAP, and global data. This data contains keys in plaintext and other sensitive information. It is critical to ensure tight security around this block.

Tamper resistance is the fourth element of a root of trust, which is essential for a hardware root of trust. Code from the outside needs to be validated prior to running it on the secure CPU. Tamper resistance can be implemented in many ways—for example, using a dedicated ROM that can be accessed only by the hardware root of trust.

The fifth element of a root of trust is a true random number generator (TRNG). This capability always produces a high level of entropy required for the various security functions. Secure, untampered access to this module is critical.

A secure clock or secure counter is the sixth element of a root of trust. It is important for applications that require a reliable time measurement. Using a secure clock is effective only if the hardware root of trust has access to a clock source that cannot be tampered with. Time is essential in a root of trust.

The seventh element of a root of trust is secure storage. Secure access to persistent storage is essential for applications requiring state knowledge. It's critical that the information cannot be tampered with, nor can the access to the information be tampered with.

The general-purpose I/O GPIO framework extension (GpioClx) simplifies the task of writing a driver for a GPIO controller device. GpioClx provides driver support for peripheral devices that connect to GPIO pins on systems.

Trusted Platform Module

Trusted Platform Module (TPM) technology is designed to provide hardware-based, security-related functions. A TPM chip is a secure crypto-processor that is designed to carry out cryptographic operations. The chip includes multiple physical security mechanisms to make it tamper resistant, and malicious software is unable to

tamper with the security functions of the TPM. Some of the key advantages of using TPM technology are that you can do the following:

Key Topic

- Generate, store, and limit the use of cryptographic keys.
- Use TPM technology for platform device authentication by using the TPM's unique RSA key, which is burned into itself.
- Help ensure platform integrity by taking and storing security measurements.

TPM is used for system integrity measurements and for key creation and use. During a system's boot process, the boot code that is loaded can be measured and recorded in the TPM. The integrity measurements can be used as evidence for how a system started and to make sure that a TPM-based key was used only when the correct software was used to boot the system.

Starting with Windows 10, the operating system automatically initializes and takes ownership of the TPM. This means that in most cases, we recommend that you avoid configuring the TPM through the TPM management console (**tpm.msc**).

NOTE Trusted Platform Modules (TPMs) and Hardware Security Modules (HSMs) provide strong hardware-based cryptographic solutions across a number of use cases, including password protection and device identification and authentication.

Key Topic

Sandboxing

Sandboxing is a strategy that isolates a test environment for applications away from critical system resources and other programs. It provides an extra layer of security that prevents malware or harmful applications from negatively affecting your system. Often virus companies launch unknown code in sandboxes to see how it reacts. Because the sandbox is isolated, it cannot affect the overall system. The sandbox provides a safe environment for opening suspicious files, running untrusted programs, or downloading URLs, without affecting the devices they are on. It can be used anytime, for any situation, to safely examine a file or code that could be malicious, before serving it up to devices—all the while keeping it isolated from a PC and the company network.

Sandboxing is used as a resource to test software that could end up being categorized as "safe" or "unsafe." As malware becomes more prevalent and dangerous, malicious applications, links, and downloads could potentially gain endless access to a network's data if they're not tested by sandbox software first. Sandboxing can be used as a tool to detect malware attacks and block them before they enter a network. The system allows your IT team to test code and understand exactly how it works

before it invades an endpoint device with malware or viruses; this gives the IT team insight and tips on what to look out for in other scenarios.

As a key measure in network and web security strategies, sandboxing provides an additional layer of security to analyze threats, separating them from the network to ensure online threats do not compromise operations. The application or file can be run if needed, with all changes being discarded after the sandbox is closed to eliminate risk of corrupted devices.

Sandbox software is available as a cloud-based or appliance-based solution and offers different advantages depending on your business needs.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 18-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 18-2 Key Topics for Chapter 18

Key Topic Element	Description	Page Number
Section	Antivirus	451
Section	Antimalware	452
List	Primary functions of an EDR system	453
Section	Data Loss Prevention	453
Section	Next-Generation Firewalls	453
Section	Host-based Intrusion Prevention System	454
Section	Host-based Intrusion Detection System	456
Section	Host-based Firewall	457
Section	Boot Integrity	458
Section	Boot Security/Unified Extensible Firmware Interface	459
List	Measured boot process	460
Section	Boot Attestation	460

Key Topic Element	Description	Page Number
Section	Database	461
Figure 18-6	Tokenization process	462
Section	Salting	462
Section	Hashing	463
Section	Application Security	463
Section	Input Validations	464
Section	Secure Cookies	465
Section	Hypertext Transfer Protocol Headers	465
Section	Code Signing	466
Section	Allow List	467
Section	Block List/Deny List	467
Figure 18-7	Security processes in the software development lifecycle (SDLC)	468
Section	Static Code Analysis	468
List	Running a static code analysis	469
Section	Manual Code Review	470
Section	Dynamic Code Analysis	470
Section	Fuzzing	471
Section	Hardening	471
Section	Open Ports and Services	471
Section	Registry	472
Section	Disk Encryption	473
List	Protection methods to harden the operating system	473
Section	Patch Management	474
Section	Self-Encrypting Drive/Full-Disk Encryption	475
Section	Hardware Root of Trust	476
List	Advantages of Trusted Platform Module (TPM) technology	478
Section	Sandboxing	478

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

endpoint protection, antivirus software, antimalware, endpoint detection and response, DLP, next-generation firewall (NGFW), host-based intrusion prevention system (HIPS), host-based intrusion detection system (HIDS), host-based firewall, boot integrity, Unified Extensible Firmware Interface (UEFI), measured boot, boot attestation, tokenization, salt, hashing, input validation, secure cookie, code signing, allow list, block listing, deny lists, static code analysis, manual code review, dynamic code analysis, fuzzing, hardening, registry, disk encryption, patch management, self-encrypting drive (SED), full-disk encryption (FDE), Opal, hardware root of trust, Trusted Platform Module (TPM), sandboxing

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What are the three strategies that antimalware software uses to protect systems from malicious software?
2. What is the first step toward achieving a trusted infrastructure on computers and networking devices?
3. In boot attestation, what is measured and committed during the boot process?
4. What places an exterior guard on the internal contents of a device?
5. What aspect of a disk array requires that replacement drives be configured to match the encryption protection at installation?



This chapter covers the following topics related to Objective 3.3 (Given a scenario, implement secure network designs) of the CompTIA Security+ SY0-601 certification exam:

- Load balancing
  - Active/active
  - Active/passive
  - Scheduling
  - Virtual IP
  - Persistence
- Network segmentation
  - Virtual local area network (VLAN)
  - Screened subnet (previously known as demilitarized zone)
  - East-west traffic
  - Extranet
  - Intranet
  - Zero Trust
- Virtual private network (VPN)
  - Always-on
  - Split tunnel vs. full tunnel
  - Remote access vs. site-to-site
  - IPsec
  - SSL/TLS
  - HTML5
  - Layer 2 tunneling protocol (L2TP)
- DNS
- Network access control (NAC)
  - Agent and agentless
- Out-of-band management
- Port security
  - Broadcast storm prevention
  - Bridge Protocol Data Unit (BPDU) guard
  - Loop prevention
  - Dynamic Host Configuration Protocol (DHCP) snooping
  - Media Access Control (MAC) filtering
- Network appliances
  - Jump servers
  - Proxy servers
    - Forward
    - Reverse
  - Network-based intrusion detection system (NIDS)/network-based intrusion prevention system (NIPS)
    - Signature-based
    - Heuristic/behavior
    - Anomaly
    - Inline vs. passive
  - HSM
  - Sensors
  - Collectors
  - Aggregators
- Firewalls
  - Web application firewall (WAF)
  - NGFW
  - Stateful
  - Stateless
  - Unified threat management (UTM)
  - Network Address Translation (NAT) gateway
  - Content/URL filter
  - Open-source vs. proprietary
  - Hardware vs. software
  - Appliance vs. host-based vs. virtual
- Access control list (ACL)
- Route security
- Quality of service (QoS)
- Implications of IPv6
- Port spanning/port mirroring
  - Port taps
- Monitoring services
- File integrity monitors

Implementing Secure Network Designs

In this chapter you learn about implementing secure network designs. We start by examining load balancing concepts and continue with an overview of network segmentation. We then take a deeper dive into the topic of virtual private networks (VPN). We touch on the topics of DNS security, network access control (NAC), and out-of-band management. The chapter continues with a discussion of port security, network appliances, access control lists (ACL), and route security and concludes by discussing quality of service (QoS), implications of IPv6, port spanning/port mirroring, monitoring services, and file integrity monitors.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 19-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 19-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Load Balancing	1
Network Segmentation	2
Virtual Private Network	3
DNS	4–5
Network Access Control	6
Out-of-Band Management	7
Port Security	8
Network Appliances	9

Foundation Topics Section	Questions
Access Control List	10
Route Security	11
Quality of Service	12
Implications of IPv6	13
Port Spanning/Port Mirroring	14
Monitoring Services	15
File Integrity Monitors	16

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. During an audit of your servers, you notice that most servers have large amounts of free disk space and have low memory utilization. Which of the following statements is correct if you migrate some of the servers to a virtual environment?
   a. You might end up spending more on licensing but less on hardware and equipment.
   b. You need to deploy load balancing and clustering.
   c. Your baselining tasks will become simpler.
   d. Servers will encounter latency and lowered throughput issues.

2. Which of the following can help defend against pivoting during a compromise?
   a. WAF
   b. Network segmentation
   c. WAP
   d. All of these answers are correct.

3. Which of the following is a true statement about VPNs?
   a. A remote-access VPN must include a host installed on the client.
   b. A clientless VPN can connect multiple concentrators together.
   c. A remote-access VPN may include a host installed on the client.
   d. A clientless VPN installs software on the host to establish the VPN connection.

4. In which type of attack does an attacker generate a high number of requests over port 53?
   a. CAM memory leak attack
   b. MITM
   c. ARP cache poisoning
   d. DNS amplification

5. In which of the following attacks does the attacker change the registration of a domain name without the permission of the original owner/registrant?
   a. Directory traversal
   b. Path traversal
   c. Dot-dot-slash
   d. Domain hijacking

6. Which type of network security control could be used to control access to a network based on the security posture?
   a. DNS
   b. NAC
   c. WAF
   d. None of these answers are correct.

7. Which type of management access would require an alternative path?
   a. Keyloggers
   b. Out-of-band management
   c. DNS
   d. NTP

8. Which feature present in most switches is used to provide access control by restricting the MAC address that can connect?
   a. DHCP snooping
   b. MITM detection
   c. ARP cache poisoning detection
   d. Port security

9. Which type of appliance secures a network by keeping machines behind it anonymous?
   a. DHCP server
   b. WLAP
   c. WAP
   d. IP proxy server

10. Which type of ACL would control access based on MAC address?
    a. Layer 3
    b. Layer 4
    c. Layer 5
    d. Layer 2

11. Which routing protocol is one of the most common ones used in route manipulation attacks?
    a. OSPF
    b. RIP
    c. IGRP
    d. BGP

12. Which QoS feature can be used for guaranteed bandwidth?
    a. WRED
    b. DSCP
    c. ToS
    d. CBWFQ

13. Which type of IPv6 address is structured like a unicast address?
    a. Multicast
    b. RIP
    c. IGRP
    d. Anycast

14. Which type of packet capture solution would be implemented as a command on a network switch?
    a. Tap
    b. Span port
    c. Sniffer
    d. None of these answers are correct.

15. Which of the following enables you, as security administrator, to identify an issue based on a change in normal behavior?
    a. Baselining
    b. Spanning
    c. Packet capture
    d. None of these answers are correct.

16. Which tool would help you identify whether the correct permissions are in place for files or folders on a system you are monitoring?
    a. NIPS
    b. NIDS
    c. FIM
    d. None of these answers are correct.

Foundation Topics



Load Balancing

Load balancing is a method used when you have redundant servers. Load-balancing clusters are multiple computers connected together for the purpose of sharing resources such as CPU, RAM, and hard disks. In this way, the cluster can share CPU power, along with other resources, and balance the CPU load among all the servers. Microsoft's Cluster Server is an example of this (although it can also act in failover mode), enabling parallel, high-performance computing. Several third-party vendors offer clustering software for operating systems and virtual operating systems as well. It is a common technique in web and FTP server farms, as well as in IRC servers, DNS servers, and NTP servers.

Active/Active

In an *active/active* scenario, each device performs work simultaneously, sharing the load. In this type of implementation, typically, two devices run the same software or service. This is an important aspect of the implementation of any type of failover or load balancing. If one device has a different version, even if it is slightly different, that can cause issues with the failover or load-balancing process.

Active/Passive

In an *active/passive* scenario, one device actively performs work while the other works in a standby nature. The passive device becomes active only when the active device goes offline or is manually switched to be the secondary.

Scheduling

In load-balancing environments, *scheduling* algorithms are used to determine which device should be used to process the incoming traffic.

Virtual IP

In load-balancing environments, a *Virtual IP* address is used to provide a single IP address for accessing the shared resource.

Persistence

Various mechanisms can be used to maintain session *persistence* in a load-balancing environment. For instance, in a load-balanced web server environment, a session cookie can be used to maintain that persistence between the client and server. Another name for session persistence is *sticky sessions*. The concept of session persistence forces a session to "stick" with the same server that it started with. In a non-persistent implementation, the session might start off with one server but end up mixed between other servers in the load-balancing environment. Persistence helps maintain efficiency of the load-balancing process.

Key Topic Network Segmentation

Firewalls can provide *network segmentation* while enforcing policies between those segments. In Figure 19-1, a firewall segments and enforces policies among three networks in the overall corporate network. The first network is the finance department, the second is the engineering department, and the third is the sales department.

Key Topic

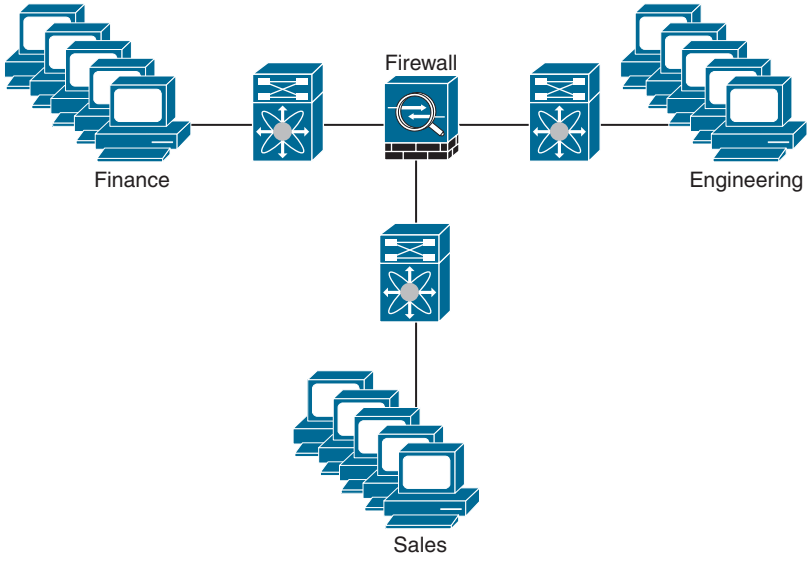


FIGURE 19-1 Firewall Providing Network Segmentation

Application-Based Segmentation and Microsegmentation

One concern to address is the machine-to-machine communication between different systems and applications. How do you also segment and protect that in an effective manner?

In today's virtualized and containerized environments, traffic between applications may never leave a physical device or server, as illustrated in Figure 19-2.

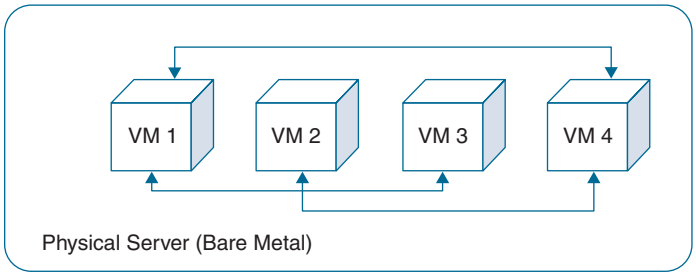


FIGURE 19-2 VM Traffic Never Leaving the Physical Server

This is why microsegmentation is so popular nowadays. A solution of the past was to include virtual firewalls between virtual machines (VMs), as shown in Figure 19-3.

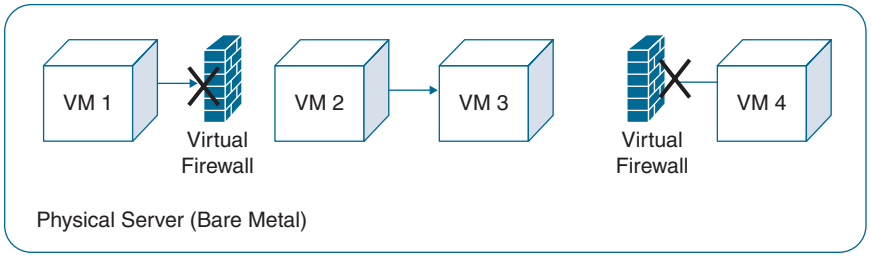


FIGURE 19-3 Virtual Firewalls for Segmentation

Machine-to-machine (or application-to-application) communication also needs to be segmented within an organization. For instance, do your Active Directory (AD) servers need to communicate to Network Time Protocol (NTP) servers? What is their relationship and data interaction?

Virtual Local Area Network

A *virtual local area network (VLAN)* is a Layer 2 broadcast domain. A careful plan of how ports or users are assigned to a specific VLAN can allow network segmentation and facilitate the implementation of access policy (for example, via network ACLs for traffic that needs to be routed across VLAN segments).

VLAN ACLs, also called VLAN maps, are not specifically Layer 2 ACLs; however, they work to limit traffic within a specific VLAN. VLAN maps can apply MAC access lists or Layer 3 and Layer 4 access lists to the inbound direction of a VLAN to provide access control.

A private VLAN can also be used to implement VLAN partitioning and control the communication among ports belonging to the same VLAN. A private VLAN includes three types of ports:

- **Promiscuous:** Devices attached to a promiscuous port can communicate with all ports within the switch, including isolated and community ports.
- **Isolated:** Devices attached to an isolated port can communicate only with the promiscuous port.
- **Community:** Devices attached to a community port can communicate with the promiscuous port and with other devices in the same community.

Screened Subnet

Firewalls can be configured to separate multiple network segments (or zones), usually now called screened subnets (previously known as demilitarized zones, or DMZs). These zones provide security to the systems that reside within them with different security levels and policies between them. Screened subnets can have several purposes; for example, they can serve as segments on which a web server farm resides or as extranet connections to a business partner. Screened subnets minimize the exposure of devices and clients on your internal network by allowing only recognized and managed services on those hosts to be accessible from the Internet. In Figure 19-4, the screened subnet hosts web servers that are accessible by internal and Internet hosts. In large organizations, you can find multiple firewalls in different segments and screened subnets.

Key Topic

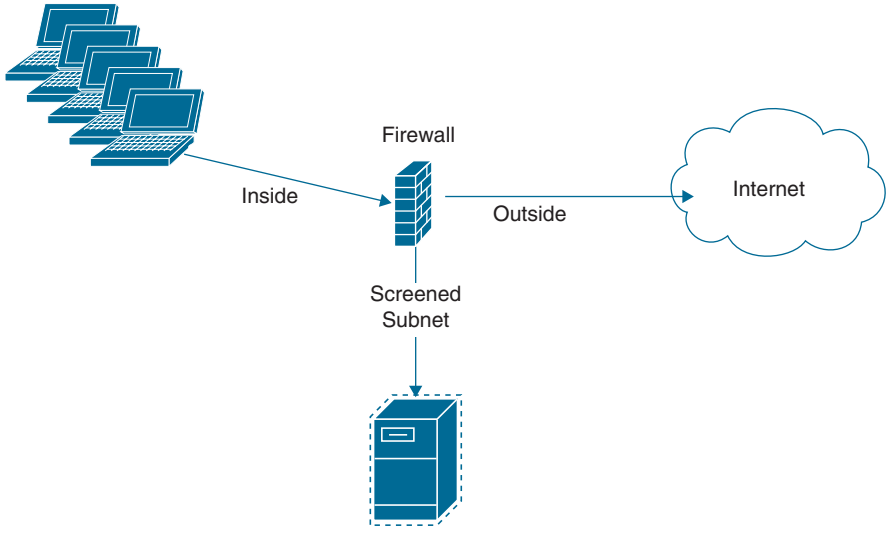


FIGURE 19-4 Screened Subnet

East-West Traffic

Firewalls can also be deployed in the data center. The placement of firewalls in the data center depends on many factors, such as how much latency the firewalls will introduce, what type of traffic you want to block and allow, and in what direction the traffic will flow (either north to south or east to west).

In the data center, traffic going from one network segment or application within the data center to another network segment or application within the data center is often referred to as east-to-west (or west-to-east) traffic. It is also known as *lateral traffic*. Figure 19-5 demonstrates *east-west traffic*.

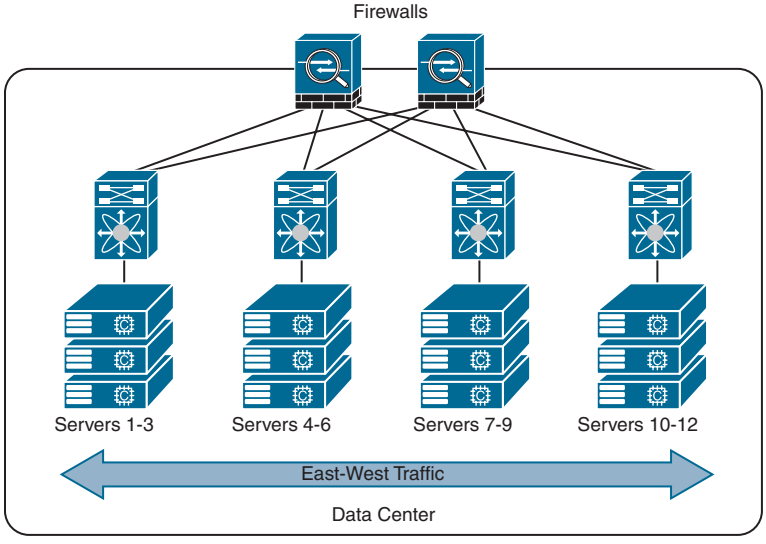


FIGURE 19-5 Data Center East-West Traffic

NOTE East-west traffic is a concept referring to network traffic flow within a data center between servers or at the access layer between devices or between network segments. North-south refers to data transfers between the LAN or data center and outside of the network, such as the Internet.

Intranets and Extranets

Intranets and *extranets* are implemented so that a company (or companies) can share its data using all the features and benefits of the Internet, while keeping that data secure within the organization, select organizations, and specific users. In the case of an intranet, only one company is involved; it could be as simple as an internal company website, or a more advanced architecture of servers, operational systems,

and networks that deploy tools, applications, and, of course, data. In the case of an extranet, multiple companies can be involved, or an organization can opt to share its data and resources with users who are not part of the organization(s). This sharing is done via the Internet, but again, is secured so that only particular people and organizations can connect.

Whether you have an intranet or an extranet, security is a major concern. Proper authentication schemes should be implemented to ensure that only the appropriate users can access data and resources. Only certain types of information should be stored on an intranet or extranet. Confidential, secret, and top-secret information should not be hosted within an intranet or extranet. Finally, the deployment of a firewall(s) should be thoroughly planned out in advance.

An example of a company that hosts an intranet and extranet is shown in Figure 19-6. Note that users from Company A can access the intranet because they work for the company. Also note that Company B can access the extranet, but not the intranet. In this example, the company (Company A) has created two (screened subnets), one for its intranet and one for its extranet. Of course, it is possible to set this up using only one screened subnet, but the access control lists on the firewall and other devices would have to be planned and monitored more carefully. If possible, separating the data into two distinct physical locations will have several benefits—namely, being more secure—although it will cost more money to do so. This all depends on the acceptable risk level of the organization and its budget!

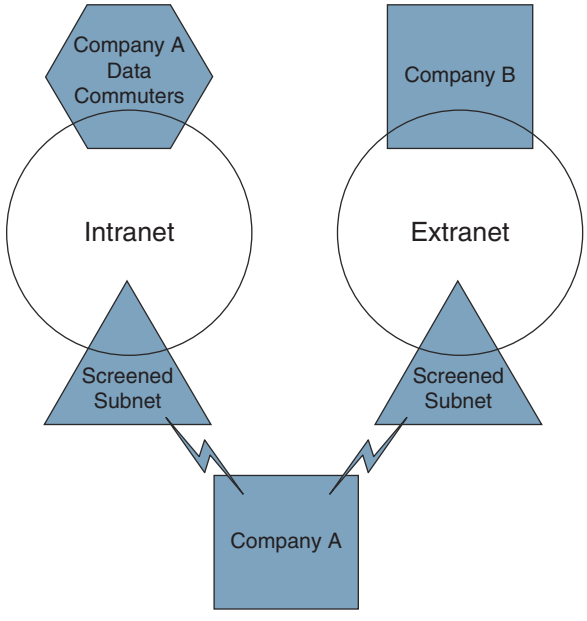


FIGURE 19-6 An Intranet and an Extranet

Zero Trust

Zero trust is a network security concept that has been around for many years. The overall principle is to not trust anything by default. Then you allow access based on various known factors that build a trust. The mechanisms for building that trust revolve around authentication, authorization, and security posture validation. Before you allow any access, these factors must be addressed.

Key Topic Virtual Private Network

A ***virtual private network (VPN)*** is a connection between two or more computers or devices not on the same private network. Generally, VPNs use the Internet to connect one host to another. It is desirable that only proper users and data sessions make their way to a VPN device; because of this, data encapsulation and encryption are used. A "tunnel' is created through any LANs and WANs that might intervene; this tunnel connects the two VPN devices together. Every time a new session is initiated, a new tunnel is created, which makes the connection secure.

Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP) are common VPN connections that can cause a lot of havoc if the security settings are not configured properly on the client side and server side.

Figure 19-7 shows a virtual private network. Note that the VPN server is on one side of the cloud, and the VPN client is on the other. The VPN client has a standard IP address to connect to its own LAN. However, it receives a second IP address from the VPN server or a DHCP device. This second IP address works "inside" the original IP address. So, the client computer has two IP addresses; in essence, the VPN address is encapsulated within the logical IP address. VPN adapters, regardless of the Internet connection used, can use MS-CHAP, as shown in Figure 19-7. To further increase authentication security, you can use a separate RADIUS server with the VPN server.

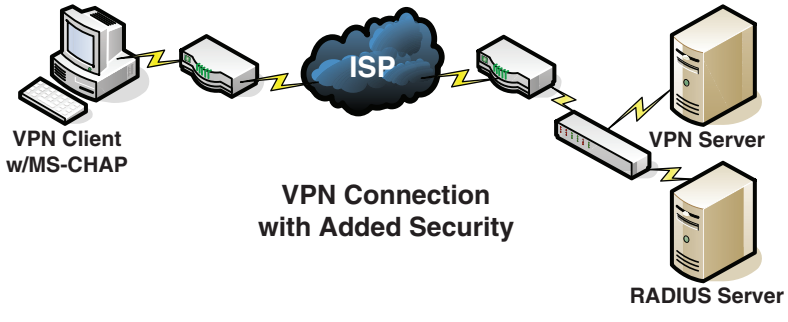


FIGURE 19-7 A VPN

Figure 19-7 illustrates a single computer connecting to a VPN server at an office, which is typical and is known as VPN remote access. However, organizations sometimes need to connect multiple offices to each other. This is done with a *site-to-site* configuration, where each site has a VPN device (SOHO router, concentrator, or server) that takes care of VPN connections for each network of computers. Site-to-site VPNs are generally more secure because you, as administrator, can specify that only specific networks can connect—and can do it in a private intranet fashion. If a company is growing, site-to-site is the way to go, whether the company is flourishing geographically or is simply inhabiting a separate space of the same building. When separate networks are connected in the same building, it is often wise to use a VPN because the physical wiring might pass through a public area.

A Microsoft VPN can be set up on a standard Windows Server by configuring Routing and Remote Access Service (RRAS). From here, remote-access policies can be created that permit or deny access to groups of users for dial-in or VPN connections. In a typical Microsoft Windows Server, you would need to set up RRAS as part of the Network Policy and Access Services role. Then you would right-click the Remote Access Logging & Policies node and access the Network Policy Server (NPS) window to create a new RRAS policy.

You don't have to use a server for incoming VPN sessions. Several vendors offer hardware appliances. Larger organizations that need hundreds of simultaneous connections should opt for a ***VPN concentrator*** as their solution. Or, it might be part of your unified threat management (UTM) solution.

Older VPNs use either PPTP (port 1723) or L2TP (port 1701) with IPsec. They can also incorporate CHAP on the client side and RADIUS servers for authentication. Newer VPNs protect traffic by using SSL or TLS (TLS being preferred as the more secure protocol). For example, OpenVPN uses this type of encryption (https://openvpn.net/). SSL/TLS solutions for VPN improve on endpoint security and enable ***always-on VPN*** functionality—where a user can always have access via the VPN without the need to periodically disconnect and reconnect.

Watch out for ***split tunneling***. This is when a client system (for example, a mobile device) can access a public network and LAN at the same time using one or more network connections. For example, a remote user might connect to the Internet through a hotel's Wi-Fi network. If the user needs to access resources on the company LAN, the VPN software takes control. But if the user needs to connect to websites on the Internet, the hotel's gateway provides those sessions. While this capability can provide for bandwidth conservation and increase efficiency, it can also bypass upper-layer security in place within the company infrastructure. While it is common, split tunneling should be tested thoroughly before you allow it. For example, you should simulate the split tunnel from a remote location and then perform vulnerability scans and capture packets. Then you can analyze the session in depth

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

and log your findings. The converse of a split tunnel configuration would be a ***full tunnel***. With a full tunnel configuration, all traffic is sent through the VPN tunnel back to the head end and out through the corporate network.

Key Topic

Remote Access vs. Site-to-Site

Typically, VPN implementations are categorized into two general groups:

- ***Site-to-site VPNs*:** Enable organizations to establish VPN tunnels between two or more network infrastructure devices in different sites so that they can communicate over a shared medium such as the Internet. Many organizations use IPsec, GRE, and MPLS VPNs as site-to-site VPN protocols.
- ***Remote-access VPNs*:** Enable users to work from remote locations such as their homes, hotels, and other premises as if they were directly connected to their corporate network.

In most cases, site-to-site VPN tunnels are terminated between two or more network infrastructure devices, whereas remote-access VPN tunnels are formed between a VPN head-end device and an end-user workstation or hardware VPN client.

Figure 19-8 illustrates a site-to-site IPsec tunnel between two sites: a site in New York (corporate headquarters) and a branch office in Raleigh, North Carolina.

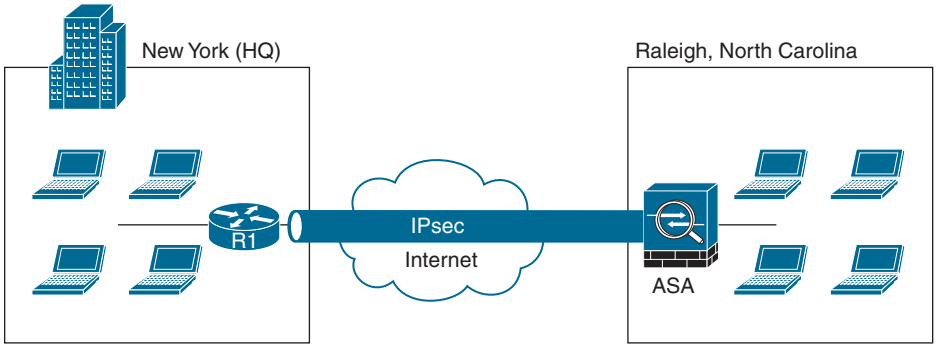


FIGURE 19-8 Site-to-Site VPN

In Figure 19-8 a Cisco router (R1) terminates an IPsec tunnel from the Cisco Firepower Threat Defense (FTD) next-generation firewall in the Raleigh office. Figure 19-9 shows a remote-access VPN.

Two clients connect to the Cisco FTD in the Raleigh office in Figure 19-9. Client 1 connects using a (TLS) VPN, and client 2 connects using IPsec.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

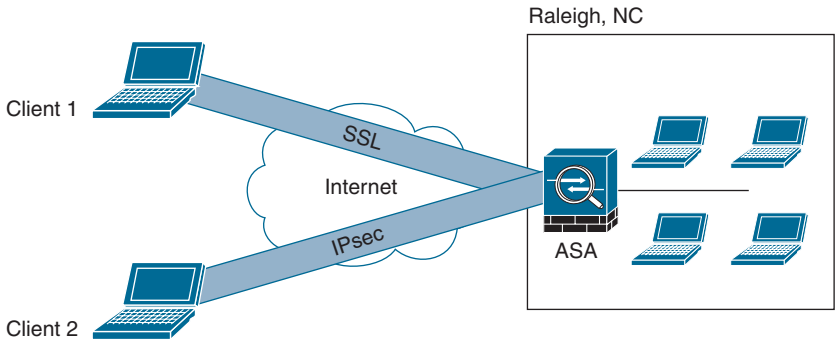


FIGURE 19-9 Remote-Access VPN

NOTE Even though SSL VPN implementations nowadays use TLS instead of SSL, most people still refer to them as "SSL VPNs." For the rest of the book, we refer to them as "SSL VPNs."

There are two main categories of remote-access VPNs:

- **Clientless:** The user connects without a client, typically using a web browser. The major benefit of clientless SSL VPNs is that you do not need a client to be installed on your PC. One of the disadvantages is that only TCP-based applications are supported. Applications using other protocols (such as UDP) cannot traverse clientless SSL VPN connections, which are typically used in kiosks, shared workstations, mobile devices, and when users just want to encrypt web traffic.
- **Client-based:** The user connects to the VPN terminating device (router, firewall, and so on) using a client. An example of a VPN client is the Cisco AnyConnect Secure Mobility Client.

IPsec

IPsec uses the Internet Key Exchange (IKE) protocol to negotiate and establish secured site-to-site or remote-access VPN tunnels. IKE is a framework provided by the Internet Security Association and Key Management Protocol (ISAKMP) and parts of two other key management protocols—namely, Oakley and Secure Key Exchange Mechanism (SKEME).

IKE is defined in RFC 2409, "The Internet Key Exchange (IKE)." IKE version 2 (IKEv2) is defined in RFC 5996, "Internet Key Exchange Protocol Version 2 (IKEv2)."

IKE has two phases. Phase 1 is used to create a secure bidirectional communication channel between the IPsec peers. This channel is known as the ISAKMP security association (SA). Phase 2 is used to negotiate the IPsec SAs.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

IKEv1 Phase 1

Within Phase 1 negotiation, several attributes are exchanged:

- Encryption algorithms
- Hashing algorithms
- Diffie-Hellman groups
- Authentication method
- Vendor-specific attributes

The following typical encryption algorithms are used in IPsec:

- Advanced Encryption Standard (AES): 128 bits long
- AES 192: 192 bits long
- AES 256: 256 bits long

TIP AES-GCM (supported only in IKEv2 implementations) stands for Advanced Encryption Standard (AES) in Galois/Counter Mode (GCM). AES-GCM is a block cipher that provides confidentiality and data-origin authentication, and greater security than traditional AES. AES-GCM supports three different key strengths: 128-, 192-, and 256-bit keys. The longer the key length, the more secure the implementation is; however, this increases the compute resources it needs for its mathematical crypto calculations. GCM is a mode of AES that is required to support the National Security Agency (NSA) Suite B. NSA Suite B is a set of cryptographic algorithms that devices must support to meet federal standards for cryptographic strength. AES-GMAC (supported only in IKEv2 IPsec proposals), which stands for Advanced Encryption Standard Galois Message Authentication Code, is a block cipher mode that provides only data-origin authentication. It is a variant of AES-GCM that allows data authentication without encrypting the data. AES-GMAC also offers the three different key strengths provided by AES-GCM (128-, 192-, and 256-bit keys).

The hashing algorithms used in IPsec include the following:

- Secure Hash Algorithm (SHA)
- Message Digest Algorithm 5 (MD5)

TIP Cisco has an excellent resource that provides an overview of all cryptographic algorithms. The same document outlines the algorithms that should be avoided and the ones that are recommended (at press time). You can access this document at https://tools.cisco.com/security/center/resources/next_generation_cryptography.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

The common authentication methods are preshared keys (where peers use a shared secret to authenticate each other) and digital certificates with the use of public key infrastructure (PKI).

Small- and medium-sized organizations use preshared keys as their authentication mechanism. Many large organizations use digital certificates for scalability, centralized management, and additional security mechanisms.

You can establish a Phase 1 SA in main mode or aggressive mode. In main mode, the IPsec peers complete a six-packet exchange in three round trips to negotiate the ISAKMP SA, whereas aggressive mode completes the SA negotiation in three packet exchanges. Main mode provides identity protection if preshared keys are used. Aggressive mode offers identity protection only if digital certificates are employed.

NOTE Cisco products that support IPsec typically use main mode for site-to-site tunnels and use aggressive mode for remote-access VPN tunnels. This is the default behavior when preshared keys are employed as the authentication method.

Figure 19-10 illustrates the six-packet exchange in main mode negotiation.

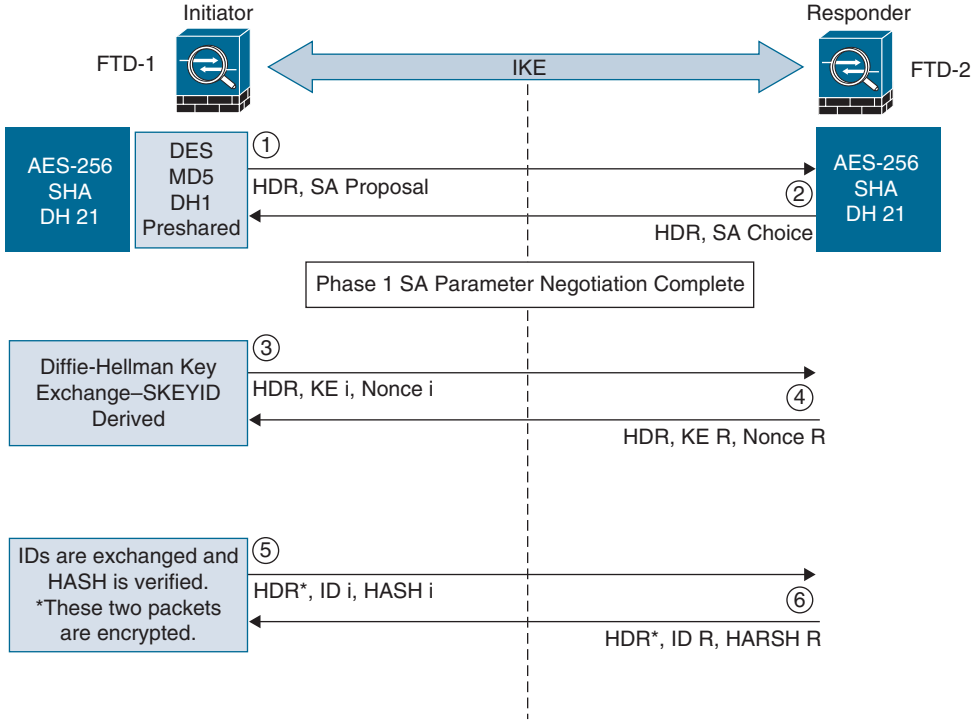


FIGURE 19-10 IPsec Phase 1 Main Mode Negotiation

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

In Figure 19-10, two firewall devices are configured to terminate a site-to-site VPN tunnel between them. The firewall labeled FTD-1 is the initiator, and FTD -2 is the responder. The following steps are illustrated in Figure 19-10:

1. FTD-1 (the initiator) has two ISAKMP proposals configured. In the first packet, FTD-1 sends its configured proposals to FTD-2.
2. FTD-2 evaluates the received proposal. Because it has a proposal that matches the offer of the initiator, FTD-2 sends the accepted proposal back to FTD-1 in the second packet.
3. The Diffie-Hellman exchange and calculation process is started. Diffie-Hellman is a key agreement protocol that enables two users or devices to authenticate each other's preshared keys without actually sending the keys over the unsecured medium. FTD-1 sends the Key Exchange (KE) payload and a randomly generated value called a *nonce*.
4. FTD-2 receives the information and reverses the equation, using the proposed Diffie-Hellman group/exchange to generate the SKEYID. The SKEYID is a string derived from secret material that is known only to the active participants in the exchange.
5. FTD-1 sends its identity information. The fifth packet is encrypted with the keying material derived from the SKEYID. The asterisk in Figure 19-10 is used to illustrate that this packet is encrypted.
6. FTD-2 validates the identity of FTD-1, and FTD-2 sends its own identity information to FTD-1. This packet is also encrypted.

TIP IKE uses UDP port 500 for communication. UDP port 500 is employed to send all the packets described in the previous steps.

For IKEv2, you can configure multiple groups. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

- **Diffie-Hellman Group 1:** 768-bit modulus. DH group 1 is considered insecure, so please do not use it.
- **Diffie-Hellman Group 2:** 1024-bit modulus. This option is no longer considered good protection.
- **Diffie-Hellman Group 5:** 1536-bit modulus. Formerly considered good protection for 128-bit keys, this option is no longer considered good protection.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

- **Diffie-Hellman Group 14:** 2048 bit modulus. Considered good protection for 192-bit keys.
- **Diffie-Hellman Group 19:** 256-bit elliptic curve.
- **Diffie-Hellman Group 20:** 384-bit elliptic curve.
- **Diffie-Hellman Group 21:** 521-bit elliptic curve.
- **Diffie-Hellman Group 24:** 2048-bit modulus and 256-bit prime order subgroup. This option is no longer recommended.

IKEv1 Phase 2

Phase 2 is used to negotiate the IPsec SAs. This phase is also known as *quick mode*. The ISAKMP SA protects the IPsec SAs because all payloads are encrypted except the ISAKMP header.

A single IPsec SA negotiation always creates two security associations—one inbound and one outbound. Each SA is assigned a unique security parameter index (SPI) value—one by the initiator and the other by the responder.

TIP The security protocols (AH and ESP) are Layer 3 protocols and do not have Layer 4 port information. If an IPsec peer is behind a PAT device, the ESP or AH packets are typically dropped because there is no Layer 4 port to be used for translation. To work around this, many vendors, including Cisco Systems, use a feature called IPsec passthrough. The PAT device that is capable of IPsec passthrough builds the Layer 4 translation table by looking at the SPI values on the packets. Many industry vendors, including Cisco Systems, implement another feature called NAT Traversal (NAT-T). With NAT-T, the VPN peers dynamically discover whether an address translation device exists between them. If they detect a NAT/PAT device, they use UDP port 4500 to encapsulate the data packets, consequently allowing the NAT device to successfully translate and forward the packets. IKEv2 enhances the IPsec interoperability between vendors by offering built-in technologies such as Dead Peer Detection (DPD), NAT Traversal (NAT-T), and Initial Contact.

Another interesting point is that if the VPN router needs to connect multiple networks over the tunnel, it must negotiate twice as many IPsec SAs. Remember, each IPsec SA is unidirectional, so if three local subnets need to go over the VPN tunnel to talk to the remote network, then six IPsec SAs are negotiated. IPsec can use quick mode to negotiate these multiple Phase 2 SAs, using the single preestablished ISAKMP (IKEv1 Phase 1) SA. The number of IPsec SAs can be reduced, however, if source and/or destination networks are summarized.

Many different IPsec attributes are negotiated in quick mode, as shown in Table 19-2.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Table 19-2 IPsec Attributes

Attribute	Possible Values
Encryption	None, AES128, AES192, AES256. AES is recommended. The higher the key length of the AES implementation, the more secure it is.
Hashing	MD5, SHA, null (SHA is recommended).
Identity information	Network, protocol, port number.
Lifetime	120–2,147,483,647 seconds. 10–2,147,483,647 kilobytes.
Mode	Tunnel or transport.
Perfect Forward Secrecy (PFS) group	None, 1, 2, or 5.

In addition to generating the keying material, quick mode also negotiates identity information. The Phase 2 identity information specifies which network, protocol, and/or port number to encrypt. Hence, the identities can vary anywhere from an entire network to a single host address, allowing a specific protocol and port.

Figure 19-11 illustrates the Phase 2 negotiation between the two routers that just completed Phase 1.

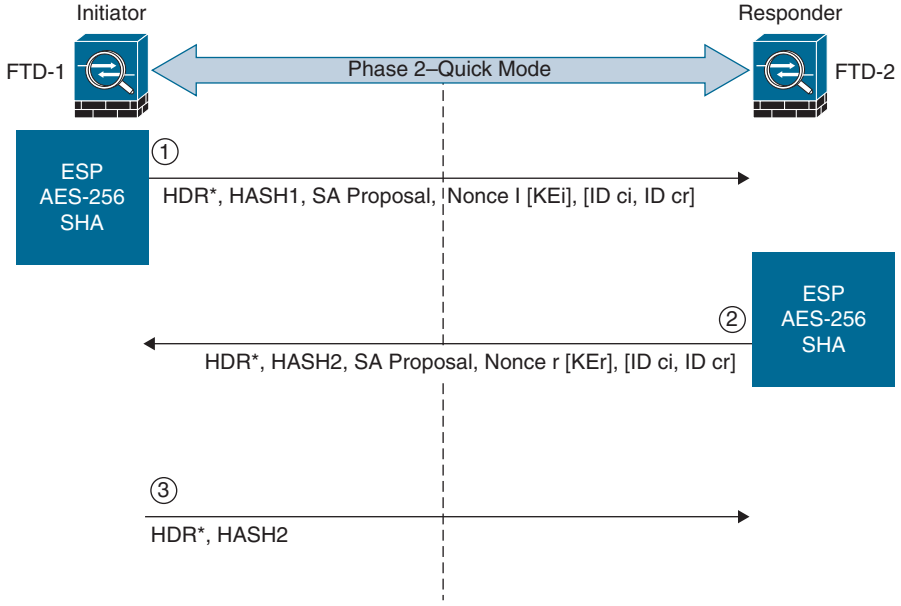


FIGURE 19-11 IPsec Phase 2 Negotiation

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

The following steps are illustrated in Figure 19-11:

1. FTD-1 sends the identity information, IPsec SA proposal, nonce payload, and (optionally) the Key Exchange (KE) payload if Perfect Forward Secrecy (PFS) is used. PFS is used to provide additional Diffie-Hellman calculations.
2. FTD-2 evaluates the received proposal against its configured proposal and sends the accepted proposal back to FTD-1, along with its identity information, nonce payload, and the optional KE payload.
3. FTD-1 evaluates the FTD-2 proposal and sends a confirmation that the IPsec SAs have been successfully negotiated. This starts the data encryption process.

IPsec uses two different protocols to encapsulate the data over a VPN tunnel:

- Encapsulation Security Payload (ESP): IP Protocol 50
- Authentication Header (AH): IP Protocol 51

ESP is defined in RFC 4303, "IP Encapsulating Security Payload (ESP)," and AH is defined in RFC 4302, "IP Authentication Header."

IPsec can use two modes with either AH or ESP:

- **Transport mode:** Protects upper-layer protocols, such as User Datagram Protocol (UDP) and TCP
- **Tunnel mode:** Protects the entire IP packet

Transport mode is used to encrypt and authenticate the data packets between the peers. A typical example is the use of GRE over an IPsec tunnel. Tunnel mode is employed to encrypt and authenticate the IP packets when they are originated by the hosts connected behind the VPN device. Tunnel mode adds an additional IP header to the packet, as illustrated in Figure 19-12.

Figure 19-12 demonstrates the major difference between transport mode and tunnel mode. It includes an example of an IP packet encapsulated in GRE and the difference when it is encrypted in transport mode versus tunnel mode. As demonstrated in Figure 19-12, tunnel mode increases the overall size of the packet in comparison to transport mode.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

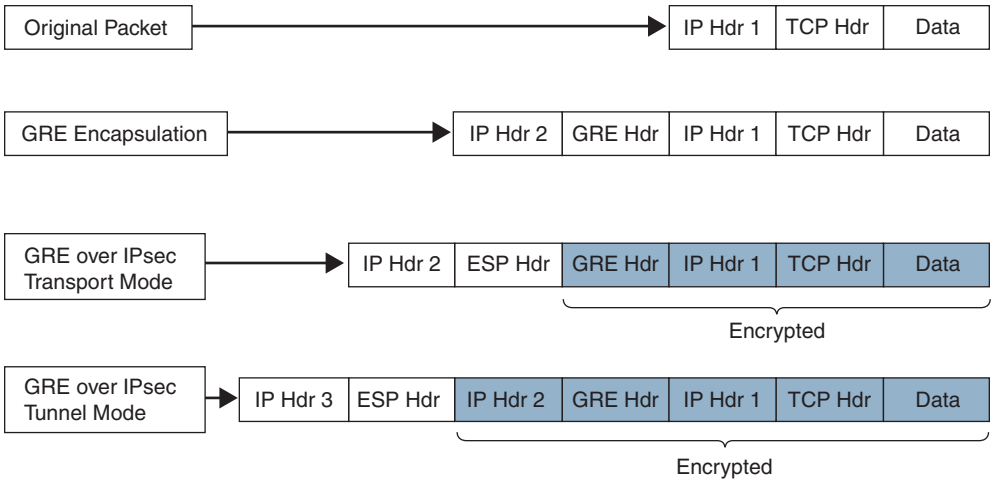


FIGURE 19-12 Transport Mode vs. Tunnel Mode in IPsec

IKEv2

IKE version 2 (IKEv2), which is defined in RFC 5996, enhances the function of performing dynamic key exchange and peer authentication. IKEv2 simplifies the key exchange flows and introduces measures to fix vulnerabilities present in IKEv1. Both IKEv1 and IKEv2 protocols operate in two phases. IKEv2 provides a simpler and more efficient exchange.

Phase 1 in IKEv2 is IKE_SA, consisting of the message pair IKE_SA_INIT. IKE_SA is comparable to IKEv1 Phase 1. The attributes of the IKE_SA phase are defined in the Key Exchange Policy. Phase 2 in IKEv2 is CHILD_SA. The first CHILD_SA is the IKE_AUTH message pair. This phase is comparable to IKEv1 Phase 2. Additional CHILD_SA message pairs can be sent for rekey and informational messages. The CHILD_SA attributes are defined in the data policy.

The following differences exist between IKEv1 and IKEv2:

- IKEv1 Phase 1 has two possible exchanges: main mode and aggressive mode. There is a single exchange of a message pair for IKEv2 IKE_SA.
- IKEv2 has a simple exchange of two message pairs for the CHILD_SA. IKEv1 uses an exchange of at least three message pairs for Phase 2. In short, IKEv2 has been designed to be more efficient than IKEv1, since fewer packets are exchanged and less bandwidth is needed compared to IKEv1.
- IKEv2 supports the use of next-generation encryption protocols and anti-DoS capabilities.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

- Despite IKEv1 supporting some of the authentication methods used in IKEv2, IKEv1 does not allow the use of Extensible Authentication Protocol (EAP). EAP enables IKEv2 to provide a solution for remote-access VPN, as well.

TIP IKEv1 and IKEv2 are incompatible protocols; consequently, you cannot configure an IKEv1 device to establish a VPN tunnel with an IKEv2 device.

Many technologies have been used for site-to-site VPN and have evolved through the years—from static traditional crypto maps to DMVPN, GETVPN, and FlexVPN.

SSL/TLS

SSL-based VPNs leverage the SSL protocol. SSL, also referred to as Transport Layer Security (TLS), is a mature protocol that has been in existence since the early 1990s. The Internet Engineering Task Force (IETF) created TLS to consolidate the different SSL vendor versions into a common and open standard.

One of the most popular features of an SSL VPN is the capability to launch a browser such as Google Chrome, Microsoft Edge, or Firefox and simply connect to the address of the VPN device, as opposed to running a separate VPN client program to establish an IPsec VPN connection. In most implementations, a clientless solution is possible. Users can access corporate intranet sites, portals, and email from almost anywhere. Even airport kiosks can establish clientless SSL VPN tunnels to access required resources. Because most people allow SSL (TCP port 443) over their firewalls, it is unnecessary to open additional ports.

The most successful application running on top of SSL is HTTP, because of the huge popularity of the World Wide Web. All the most popular web browsers in use today support HTTP over SSL/TLS (HTTPS). This ubiquity, if used in remote-access VPNs, provides some appealing properties:

- **Secure communication using cryptographic algorithms:** HTTPS/TLS offers confidentiality, integrity, and authentication.
- **Ubiquity:** The ubiquity of SSL/TLS makes it possible for VPN users to remotely access corporate resources from anywhere, using any PC, without having to preinstall a remote-access VPN client.
- **Low management cost:** The clientless access makes this type of remote-access VPN free of deployment costs and free of maintenance problems at the end-user side. This is a huge benefit for IT management personnel, who

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

would otherwise spend considerable resources to deploy and maintain their remote-access VPN solutions.

- **Effective operation with a firewall and NAT:** SSL VPN operates on the same port as HTTPS (TCP/443). Most Internet firewalls, proxy servers, and NAT devices have been configured to correctly handle TCP/443 traffic. Consequently, there is no need for any special consideration to transport SSL VPN traffic over the networks. This has been viewed as a significant advantage over native IPsec VPN, which operates over IP protocol 50 (ESP) or 51 (AH), which in many cases needs special configuration on the firewall or NAT devices to let traffic pass through.

As SSL VPN evolves to fulfill another important requirement of remote-access VPNs (namely, the requirement of supporting any application), some of these properties are no longer applicable, depending on which SSL VPN technology the VPN users choose. But overall, these properties are the main drivers for the popularity of SSL VPNs in recent years and are heavily marketed by SSL VPN vendors as the main reasons for IPsec replacement.

NOTE OpenVPN is an example of an SSL/TLS-based VPN open-source implementation. You can download and obtain detailed information about OpenVPN at https://openvpn.net.

Today's SSL VPN technology uses SSL/TLS for secure transport and employs a heterogeneous collection of remote-access technologies such as reverse proxy, tunneling, and terminal services to provide users with different types of access methods that fit different environments. Subsequent chapters examine some commonly used SSL VPN technologies, such as the following:

- Reverse proxy technology
- Port-forwarding technology and smart tunnels
- SSL VPN tunnel client (AnyConnect Secure Mobility Client)
- Integrated terminal services

HTTPS provides secure web communication between a browser and a web server that supports the HTTPS protocol. SSL VPN extends this model to allow VPN users to access corporate internal web applications and other corporate application servers that might or might not support HTTPS, or even HTTP. SSL VPN does this by using several techniques that are collectively called *reverse proxy technology*.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

A reverse proxy is a proxy server that resides in front of the application servers (normally web servers) and functions as an entry point for Internet users who want to access the corporate internal web application resources. To external clients, a reverse proxy server appears to be the true web server. Upon receiving the user's web request, a reverse proxy relays the user request to the internal web server to fetch the content on behalf of the user and then relays the web content to the user with or without presenting additional modifications to the data.

Many web server implementations support reverse proxy. One example is the mod_proxy module in Apache. With so many implementations, you might wonder why you need an SSL VPN solution to have this functionality. The answer is that SSL VPN offers much more functionality than traditional reverse proxy technologies:

- SSL VPN can transform complicated web and some non-web applications that simple reverse proxy servers cannot handle. The content transformation process is sometimes called *webification*. For example, SSL VPN solutions enable users to access Windows or Linux file systems. The SSL VPN gateway must be able to communicate with internal Windows or Linux servers and "webify" the file access in a web browser–presentable format for the VPN users.
- SSL VPN supports a wide range of business applications. For applications that cannot be webified, SSL VPN can use other resource access methods to support them. For users who demand ultimate access, SSL VPN provides network-layer access to directly connect a remote system to the corporate network, in the same manner as an IPsec VPN.
- SSL VPN provides a true remote-access VPN package, including user authentication, resource access privilege management, logging and accounting, endpoint security, and user experience.

The reverse proxy mode in SSL VPN is also known as *clientless web access* or just *clientless access* because it does not require any client-side applications to be installed on the client machine. Client-based SSL VPN provides a solution where you can connect to the corporate network by just pointing your web browser to the Cisco ASA without the need of additional software being installed on your system.

An example of an SSL/TLS VPN used in the industry today is the Cisco ASA. In the current software release, Cisco ASA supports all three flavors of SSL/TLS VPN:

- **Clientless:** In the clientless mode, the remote client needs only an SSL-enabled browser to access resources on the private network of the security appliances. SSL clients can access internal resources such as web servers, application servers, and even Windows file shares over the SSL/TLS tunnel.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

- **Thin client:** In the thin client mode, the remote client needs to install a small Java-based applet to establish a secure connection to the TCP-based internal resources. SSL clients can access internal resources such as HTTP, HTTPS, and SSH servers.
- **Full tunnel:** In the full tunnel client mode, the remote client needs to install an SSL VPN client first that can give full access to the internal private network over an SSL tunnel. Using the full tunnel client mode, remote machines can send all IP unicast traffic such as TCP-, UDP-, or even ICMP-based traffic. SSL clients can access internal resources such as HTTP, HTTPS, DNS, and SSH servers.

In some documentation, clientless and thin client solutions are grouped under one umbrella and classified as *clientless SSL VPNs*.

HTML5

With the advent of *HTML5*, a new type of VPN has now become popular. That is the use of HTML5 as a VPN portal. This allows the user to have only a browser as the VPN client. A number of vendors out there today offer this as a solution. An HTML5-based remote-access solution doesn't require browser plug-ins (or worrying about constantly updating insecure plug-ins) and is considered a secure remote-access solution alternative to using an SSL- or TLS-based VPN. This easy-to-use pure browser-based solution can increase productivity as well as security.

Layer 2 Tunneling Protocol

The *Layer 2 Tunneling Protocol (L2TP)* is used to connect VPNs. In essence, it creates an unencrypted tunnel if used by itself (which would be unwise). It does not include confidentiality or encryption on its own, but when paired with a security protocol such as IPsec, it is considered a formidable tunneling protocol.

Its starting point is based on the Layer 2 Forwarding Protocol (L2F) and PPTP. The latest version is L2TPv3, which has improved encapsulation and increased security features. Servers and other devices accepting incoming VPN connections need to have inbound port 1701 open.

When installed on a Windows Server, L2TP uses a PKI. Valid certificates need to be downloaded to clients before they can make a VPN connection to the server. Security must be configured on the server side and client side. Generally, the IPsec protocol is used to accomplish the secure connection within the L2TP tunnel.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

DNS

DNS security in an organization is an essential layer of security that is fairly easy to address. The first thing that you need to address when it comes to DNS traffic in your organization is where it is going. What DNS servers are your users utilizing? Getting a handle on this is the first step to securely implementing DNS security. If your users are able to use any DNS server out on the Internet, this opens them up to a number of attacks. So, first, you need to have visibility on your network into what DNS servers are being used. From there, you should implement a secure DNS solution and put procedures, policies, and controls in place to force the use of only those DNS servers that you approve. Many organizations utilize DNS as a base layer of security, which includes web content control and malware protection. Blocking DNS requests to malicious sites is a quick and efficient way to leverage DNS to control user access to malware. This type of service is typically offered by third-party vendors. A great example of this type of service is Cisco OpenDNS.

Cisco acquired OpenDNS, which provides DNS services, threat intelligence, and threat enforcement at the DNS layer. OpenDNS has a global network that delivers advanced security solutions (as a cloud-based service), regardless of where the customer offices or employees are located. This service is extremely easy to deploy and easy to manage. It also incorporates the innovative advancements to threat research and threat-centric security that OpenDNS has developed to block advanced cybersecurity threats with other security and networking products. Millions of people use OpenDNS, including thousands of companies, from Fortune 500 enterprises to small businesses.

OpenDNS provides a free DNS service for individuals, students, and small businesses. You can just configure your endpoint (laptop, desktop, mobile device, server, or your DHCP server) to point to OpenDNS servers: 208.67.222.222 and/or 208.67.220.220.

OpenDNS also provides the following premium services:

- **Cisco Umbrella:** This enterprise advanced network security service protects any device, anywhere. This service blocks known malicious sites from being "resolved" in DNS. It provides an up-to-the-minute view and analysis of at least 2 percent of the world's Internet activity to stay ahead of attacks. This service provides threat intelligence by seeing where attacks are being staged on the Internet.
- **Investigate:** This premium service provides information on where attacks are forming, enabling you to investigate incidents faster and prioritize them better. With the Investigate service, you can see up-to-the-minute threat data and historical context about all domains on the Internet and respond quickly to

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

critical incidents. It provides a dynamic search engine and a RESTful API that you can use to automatically bring critical data into the security management and threat intelligence systems deployed in your organization. It also provides predictive threat intelligence using statistical models for real-time and historical data to predict domains that are likely malicious and could be part of future attacks.

Key Topic Network Access Control

Throughout this book, we have mentioned several types of networking technologies and design elements. But whichever you choose to use, it needs to be controlled in a secure fashion. *Network access control (NAC)* does this by setting the rules by which connections to a network are governed. Computers attempting to connect to a network are denied access unless they comply with rules pertaining to levels of antivirus/antimalware protection, system updates, and so on…effectively weeding out those who would perpetuate malicious attacks. The client computer continues to be denied until it has been properly updated, which in some cases can be taken care of automatically by the NAC solution. This often requires some kind of preinstalled software (an agent) on the client computer, or the computer is scanned by the NAC solution remotely—which would be known as agentless. When agents are used, there is usually a posture or health check at the host or the endpoint. Often this is tied into role-based access. There are two types of agents: persistent and dissolvable. Persistent agents are installed on the target device and can be used over and over again. Dissolvable agents provide for one-time authentication and are then deleted. Agentless NAC systems are also available, but they offer less control and fewer inspection capabilities.

Some companies (such as Cisco) offer hardware-based NAC solutions, whereas other organizations offer paid software-based NAC solutions and free ones such as PacketFence (https://packetfence.org), which is open source.

The IEEE 802.1X standard, known as port-based network access control, or PNAC, is a basic form of NAC that enables the establishment of authenticated point-to-point connections, but NAC has grown to include software; 802.1X is now considered a subset of NAC.

Key Topic Out-of-Band Management

Historically, management tools have been in-band, meaning that an administrator connects locally through the main company network itself. In-band management is very common. However, you might need to take an alternate path to manage

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

network devices. In this case, you require out-of-band management. This approach is common for devices that do not have a direct network connection, such as UPSs, and environmental controls. More importantly, in a mission-critical network, SNMP-based in-band management tools are not enough, because if the main LAN fails, there is no path to manage the devices. That's when out-of-band management becomes crucial. *Out-of-band* tools that are based on serial access, USB, or even an entirely separate LAN are required if the main network fails. The key here is that an access switch of some kind that is independent from the main (in-band) network is available to the managed devices. The decision of whether to use in-band or out-of-band management (or both) should be weighed carefully. Remember that every additional system or network increases the security risk and needs to be protected appropriately—on a physical and a logical level. Then, this needs to be balanced against the level of criticality of your systems. Typically, in today's networks, an out-of-band network or management network is built for the purposes of providing an alternate communication path for managing network devices.

Key Topic Port Security

Port security is a security feature present in the routers and switches from most respected vendors, and it is used to provide access control by restricting the Media Access Control (MAC) addresses that can be connected to a given port. This differs from a MAC access list because it works only on the source MAC address without matching the MAC destination.

Port security works by defining a pool of MAC addresses that are allowed to transmit on a device port. The pool can be statically defined or dynamically learned. Compared to a MAC access list, which would need to be implemented on each port and have static entries, the dynamically learned method reduces the administrative overhead related to the port access control implementation.

When a frame is received on the port, the port security feature checks the source MAC address of the frame. If it matches an allowed MAC address, the frame is forwarded; otherwise, the frame is dropped.

In addition to the drop frame coming from an unauthorized MAC address, port security raises a security violation. A security violation is raised under the following circumstances:

- If a MAC address that is configured or dynamically learned on one port is seen on a different port in the same VLAN. This is referred to as a *MAC move*.
- If the maximum number of MAC addresses allowed for a port is reached and the incoming MAC is different from the one already learned.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Broadcast Storm Prevention

A broadcast storm happens when an overwhelming amount of broadcast traffic hits a network segment. This leads to the overconsumption of network resources and can lead to a denial-of-service (DoS) situation affecting the availability of critical corporate resources. Here, we describe some Layer 2 security best practices for securing your infrastructure.

Key Topic

Bridge Protocol Data Unit Guard

With Bridge Protocol Data Unit (BPDU) guard, if BPDUs show up where they should not, the switch protects itself. BPDU guard is configured on a network switch to prevent the switch from receiving STP BPDUs. Example 19-1 demonstrates a configuration for a Cisco IOS device.

Example 19-1 Configuration of BPDU Guard

```
CatSwitch-IOS(config)# spanning-tree portfast bpduguard
CatSwitch-IOS(config)
```

Loop Prevention

Some switches come with the ability to enable *loop protection* within the firmware. For example, you should enable the Spanning Tree Protocol (STP)—that is, if it isn't enabled by default. STP—as well as Rapid STP and multiple STP—builds a logical loop-free topology for the Ethernet network and can recognize and bypass improper connections. However, it's preferable to prevent the problem from happening in the first place. So, remember that the switch needs to be physically secured, most likely in a server room with some type of access control system, and all switches should be checked for internal cable loops. It sounds so simple, but it is commonly overlooked by many companies. Also, you should disable any unused ports on the switch, if the switch has that capability. Finally, you should employ good cable management: label ports and cables and organize patch cables as necessary.

Key Topic

Dynamic Host Configuration Protocol Snooping

DHCP snooping prevents rogue DHCP servers from impacting the network. It is implemented to address DHCP security on your network. It is configured on your network switches. The idea behind DHCP snooping is that you can determine if an ARP packet is valid based on IP-to-MAC address bindings stored in a trusted database. This database is called the "DHCP snooping binding database."

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

In MAC spoofing, an attacker masks the MAC address of his or her computer's network adapter with another number. This can allow a user to disguise a computer as another system on or off the network, thus fooling a switch and possibly gaining access. The port security feature found on switches and NAC systems or other application layer solutions can help to prevent this issue, but in the case of an insider attempting it, a router should also be configured to accept only a certain number of static MAC addresses. This attack can be enhanced by initiating a *DHCP starvation attack*, which works by broadcasting DHCP requests with spoofed MAC addresses, which can ultimately exhaust the address space available to DHCP servers. To help prevent this, you should enable DHCP snooping. Also, you need to close off any untrusted switch interfaces—meaning ones that connect outside the network or firewall. Finally, another related attack is *ARP spoofing*, which occurs when an attacker can make a system appear as the destination host sought by the sender, with obvious repercussions. You can prevent this type of attack by reducing the time an entry stays in the ARP cache—on the switch and on the clients; checking/removing static ARP entries; using dynamic ARP inspection; and also using the aforementioned DHCP snooping.

Media Access Control Filtering

Many wireless access points (WAPs) come with a built-in firewall. If the firewall is utilized, the stateful packet inspection (SPI) option and NAT filtering should be enabled. The WAP might also have the capability to be configured for *MAC filtering* (a basic form of network access control), which can filter out which computers can access the wireless network. The WAP does this by consulting a list of MAC addresses that have been previously entered. Only the network adapters with those corresponding MAC addresses can connect; no one else can join the wireless network. In some cases, a device might broadcast this MAC table. If this is the case, look for an update for the firmware of the WAP, and again, attempt to fine-tune the broadcast range of the device so that it does not leak out to other organizations. Because MAC filtering and a disabled SSID can be easily circumvented using a network sniffer, it is very important to also use strong encryption and possibly consider other types of network access control (such as 802.1X) and external authentication methods (such as RADIUS).

Key Topic Network Appliances

The following sections cover network appliances that you might find in your environment, including jump servers and proxy servers. Additionally, we cover network-based intrusion detection and network-based intrusion prevention systems. We discuss the topics related to NIPS, such as signature based, heuristic/behavior, and anomaly detection while also touching on inline versus passive implementation.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

We then dive into the concept of HSM, sensors, collectors, and aggregators. We conclude with a discussion of firewalls, which include the topics of web application firewall, next-generation firewall, stateful, stateless, unified threat management, Network Address Translation gateway, content/URL filter, open-source versus proprietary, hardware versus software, and appliance versus host-based versus virtual.

Jump Servers

A jump server, also known as a "jump box," is a device on your network that is typically used to access other devices on the network. For instance, you might have direct access to a Windows system via remote desktop protocol from outside your network or on a secure enclave. Gaining access to that machine then allows you to access devices that you do not have a direct access to. An example would be a jump server installed on a screened subnet for managing the servers within the screened subnet.

Proxy Servers

A *proxy server* acts as an intermediary for clients, usually located on a LAN, and the servers that they want to access, usually located on the Internet. By definition, *proxy* means go-between, or mediator, with the server acting as such a mediator in between a private network and a public network. The proxy server evaluates requests from clients and, if they meet certain criteria, forwards them to the appropriate server. There are several types of proxies, including a couple you should know for the exam:

- **IP proxy:** Secures a network by keeping machines behind it anonymous; it does this through the use of NAT. For example, a basic four-port router can act as an IP proxy for the clients on the LAN it protects. An IP proxy can be the victim of many network attacks, especially DoS attacks. Regardless of whether the IP proxy is an appliance or a computer, it should be updated regularly, and its log files should be monitored periodically and audited according to organization policies.
- **Caching proxy:** Attempts to serve client requests without actually contacting the remote server. Although there are FTP and SMTP proxies, among others, the most common caching proxy is the HTTP proxy, also known as a web proxy, which caches web pages from servers on the Internet for a set amount of time. Examples of caching proxies include WinGate (for Windows systems) and SquidProxies (commonly used on Linux-based systems). Figure 19-13 shows a caching proxy.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

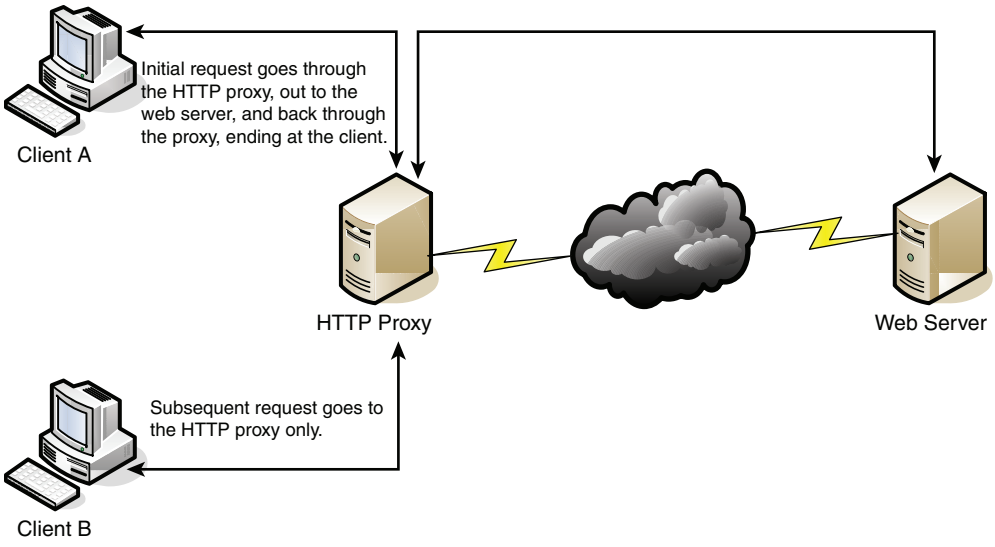


FIGURE 19-13 An HTTP Proxy in Action

For example, let's say your coworker (Client A) accessed www.google.com and was the first person to do so on the network. This client request goes through the HTTP proxy and is redirected to Google's web server. As the data for Google's home page comes in, the HTTP proxy stores or caches that information. When another person on your network (Client B) makes a subsequent request for www.google.com, the bulk of that information comes from the HTTP proxy instead of from Google's web server. This is done to save bandwidth on the company's Internet connection and to increase the speed at which client requests are carried out.

Most HTTP proxies check websites to verify that nothing has changed since the last request. Because information changes quickly on the Internet, a time limit of 24 hours is common for storing cached information before it is deleted. Web browsers make use of a proxy autoconfiguration (PAC) file, which defines how the browser can automatically choose a proxy server. The file itself and the embedded JavaScript function pose a security risk in that the file can be exploited and modified, ultimately redirecting the user to unwanted (and potentially malicious) websites. You should consider disabling PAC files and autoconfiguration in general within client web browsers.

Other types of proxies are available to apply policies, block undesirable websites, audit employee usage, and scan for malware. One device or computer might do all these things or just one or two. It depends on the software used or appliance installed. Reverse proxies can also be implemented to protect a DMZ server's identity or to provide authentication and other secure tasks. This is done when users

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

on the Internet are accessing server resources on your network. Generally, a proxy server has more than one network adapter so that it can connect to the various networks for which it is acting as a mediator. Each of the network adapters in a proxy should be periodically monitored for improper traffic and for possible network attacks and other vulnerabilities. A proxy server might be the same device as a firewall, or it could be separate. Because of this, a multitude of network configurations are possible. Proxy servers, especially HTTP proxies, can be used maliciously to record traffic sent through them; because most of the traffic is sent in unencrypted form, this could be a security risk. A possible mitigation for this is to chain multiple proxies together in an attempt to confuse any onlookers and potential attackers.

Most often, a proxy server is implemented as a ***forward proxy***. This means that clients looking for websites, or files via an FTP connection, pass their requests through to the proxy. However, there is also a ***reverse proxy***, where multiple HTTP or FTP servers use a proxy server and send out content to one or more clients. These HTTP and FTP servers could be located in a server farm or similar grouping, and the reverse proxy might also undertake the role of load balancer in this situation. A reverse proxy can act as another layer of defense for an organization's FTP or HTTP servers. An application proxy might be used as a reverse proxy; for example, Microsoft's Web Application Proxy, which enables remote users to connect to the organization's internal network to access multiple servers. These are often multipurpose by design, allowing HTTP, FTP, email, and other types of data connections. However, you could have a single application stored on several servers. Those servers can work together utilizing clustering technology. The clustering might be controlled by the servers themselves, or more commonly, a load balancer can be installed in front of the servers that distributes the network load among them. That load balancer in effect acts as a reverse proxy.

Regardless of the type of proxy used, it often modifies the requests of the "client computer," whatever that client is, providing for a level of anonymity. But in some cases, you might need a proxy that does not modify requests. This is known as a *transparent proxy*. Although it allows for increased efficiency, there is less protection for the client system.

Network-Based Intrusion Detection System/Network-Based Intrusion Prevention System

It is critical that you should be able to differentiate between a network-based intrusion *detection* system (NIDS) and a network-based intrusion *prevention* system (NIPS) for the exam. Although a great many attacks can hamper an individual computer, just as many network attacks could possibly take down a server, switch, router, or even an entire network. Network-based intrusion detection systems were

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

developed to detect these malicious network attacks, and network-based intrusion prevention systems were developed in an attempt to prevent them.

NIDS

A ***network-based intrusion detection system (NIDS)***, by definition, is a type of IDS that attempts to detect malicious network activities—for example, port scans and DoS attacks—by constantly monitoring network traffic. It can also be instrumental in rogue machine detection, including rogue desktops, laptops, and mobile devices, as well as rogue access points, DHCP servers, and network sniffers. Examples of NIDS solutions include open-source products such as Snort (www.snort.org/), Zeek (www.bro.org/), and many other commercial hardware- and software-based products. A NIDS should be situated at the entrance or gateway to your network. It is not a firewall but should be used with a firewall. Because the NIDS inspects every packet that traverses your network, it needs to be fast; basically, the slower the NIDS, the slower the network. So, the solution itself, the computer or device it is installed on, and the network connections of that computer or device all need to be planned out accordingly to ensure that the NIDS does not cause network performance degradation.

Figure 19-14 illustrates how a NIDS might be implemented on a network. Often it is placed in front of a firewall. The NIDS detects attacks and anomalies and alerts you if they occur, whereas the firewall does its best to prevent those attacks from entering the network. However, a NIDS could be placed behind the firewall, or you might have multiple NIDS points strategically placed around the network. If the NIDS is placed in front of the firewall, it generates a lot more administrator alerts, but they can usually be whittled down within the firmware or software of the device running the NIDS. Regardless of where the NIDS is located, a network administrator should monitor traffic from time to time; to do so, the computer, server, or appliance that has the NIDS installed should have a network adapter configured to work in ***promiscuous mode***. This passes all traffic to the CPU, not just the frames addressed to it.

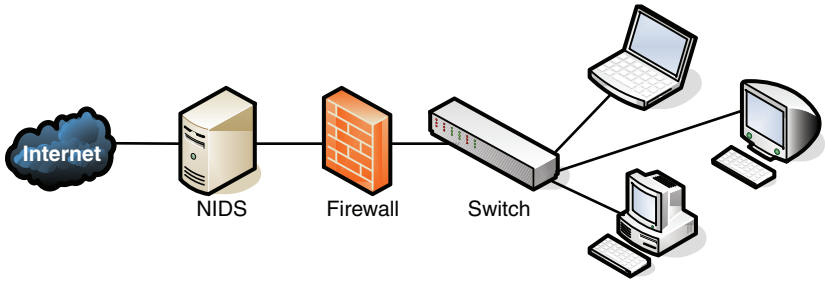


FIGURE 19-14 NIDS Placement in a Network

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

The beauty of a NIDS is that you might get away with one or two NIDS points on the network, and do away with some or all the host-based intrusion detection systems (HIDS) installed on individual computers, effectively lowering the bottom line while still doing a decent job of mitigating risk. A couple of disadvantages of a NIDS, aside from possible network performance issues, are that it might not be able to read encrypted packets of information and will not detect problems that occur on an individual computer. Therefore, to secure a network and its hosts, many organizations implement a mixture of NIDS and HIDS. If a NIDS is placed in front of the firewall, it is subject to attack; therefore, it should be monitored and updated regularly. Some NIDS solutions auto-update. Finally, the biggest disadvantage of a NIDS is that it is passive, meaning it only *detects* attacks; to protect against, or *prevent*, these attacks, you need something active. In this case, you need a NIPS.

NIPS

A ***network-based intrusion prevention system (NIPS)*** is designed to inspect traffic and, based on its configuration or security policy, either remove, detain, or redirect malicious traffic that it becomes aware of. The NIPS (as well as the NIDS) is considered to be an *application-aware device*, meaning it can divine different types of packets, define what application they are based on, and ultimately permit or disallow that traffic on the network. More and more companies are offering NIPS solutions in addition to, or instead of, NIDS solutions. Examples of NIPS solutions include Check Point security appliances (www.checkpoint.com), and the aforementioned Snort, which is actually a NIDS/NIPS software package that should be installed on a dual-homed or multihomed server. Not only can a NIPS go above and beyond a NIDS by removing or redirecting malicious traffic, but it can also redirect a recognized attacker to a single computer known as a padded cell, which contains no information of value and has no way out.

Like a NIDS, a NIPS should sit inline on the network, often in front of the firewall, although it could be placed elsewhere, depending on the network segment it protects and the network architecture. Whereas many NIPS solutions have two connections only and are known as perimeter solutions, other NIPS appliances have up to 16 ports enabling many points of detection on the network; these are known as network "core" devices. Regardless of the solution you select, as packets pass through the device, they are inspected for possible attacks. These devices need to be accurate and updated often (hopefully automatically) to avoid the misidentification of legitimate traffic, or worse, the misidentification of attacks. If the NIPS blocks legitimate traffic, this misidentification would be known as a ***false positive***, and effectively could deny service to legitimate customers, creating a self-inflicted denial of service of sorts.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

If the NIPS does not have a particular attack's signature in its database and lets that attack through thinking it is legitimate traffic, it is known as a ***false negative***, also bad for obvious reasons! Many NIPS systems can monitor for attack signatures and anomalies.

Another type of error that can occur with NIDS and NIPS is a subversion error; this occurs when the NIDS/NIPS has been altered by an attacker to allow for false negatives, ultimately leading to attacks creeping into the network. This situation can be deadly because the NIDS/NIPS often is the first point of resistance in the network. To protect against this type of error, some devices have the capability to hide or mask their IP address. They might also come with an internal firewall. It is also important to select a NIPS solution that has a secure channel for the management console interface.

One advantage of newer NIPS solutions is that some of them can act as protocol analyzers by reading encrypted traffic and stopping encrypted attacks. In general, the beauty of a NIPS is that it can protect non-computer-based network devices such as switches, routers, and firewalls. However, the NIPS is considered a single point of failure because it sits inline on the network. As a result, some organizations opt to install a bypass switch, which also enables the NIPS to be taken offline when maintenance needs to be done.

A vital NIPS consideration is whether to implement a fail-close or fail-open policy—in essence, deciding what will happen if the NIPS fails. Fail-close means that all data transfer is stopped, while fail-open means that data transfers (including potential attacks) are passed through. Let's consider an example. Say that the NIPS was protecting an individual server (or router) and had a certain level of control over that system. Now let's say that the NIPS failed. In a fail-close scenario, it would disconnect the system that it is protecting, stopping all data transfer. This outcome is unacceptable to some organizations that require near 100 percent uptime. These organizations are willing to accept additional risk and therefore are more receptive to a fail-open scenario. However, in this case, if the NIPS fails, it continues to pass all traffic to the "protected" system, which could include possible attacks. Sometimes, fail-open scenarios are necessary. In these cases, defense in depth is the best strategy. For instance, you might opt to have a firewall filter the bulk of traffic coming into the network but have the NIPS filter only specific traffic, reducing the chances of NIPS failure. This layered approach can offer greater security with less chance of attacks passing through, but often comes with increased cost and administration.

Summary of NIDS vs. NIPS

Table 19-3 summarizes the advantages and disadvantages of NIDS versus NIPS.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Table 19-3 NIDS vs. NIPS

Type of System	Summary	Disadvantage/Advantage	Example
NIDS	Detects malicious network activities	Pro: Only a limited number of NIDSs are necessary on a network. Con: Only detects malicious activities.	Snort Bro IDS
NIPS	Detects, removes, detains, and redirects traffic	Pro: Detects and mitigates malicious activity. Pro: Can act as a protocol analyzer. Con: Uses more resources. Con: Possibility of false positives and false negatives.	Check Point Systems solutions

Signature-Based

Pattern matching is a methodology in which the intrusion detection device searches for a fixed sequence of bytes within the packets traversing the network. Generally, the pattern is aligned with a packet that is related to a specific service or, in particular, associated with a source and destination port. This approach reduces the amount of inspection made on every packet. However, it is limited to services and protocols that are associated with well-defined ports. Protocols that do not use any Layer 4 port information are not categorized. Examples of these protocols are Encapsulated Security Payload (ESP), Authentication Header (AH), and Generic Routing Encapsulation (GRE).

This tactic uses the concept of signatures. A *signature* is a set of conditions that point out some type of intrusion occurrence. For example, if a specific TCP packet has a destination port of 1234 and its payload contains the string ff11ff22, a signature can be configured to detect that string and generate an alert.

Alternatively, the signature could include an explicit starting point and endpoint for inspection within the specific packet.

Here are some of the benefits of the plain pattern-matching technique:

- Direct correlation of an exploit
- Trigger alerts on the pattern specified
- The capability to apply across different services and protocols

One of the main disadvantages is that pattern matching can lead to a considerably high rate of *false positives* or *false negatives*, which were previously mentioned.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

To address some of these limitations, a more refined method was created. This methodology is called *stateful pattern-matching recognition*. This process dictates that systems performing this type of signature analysis must consider the chronological order of packets in a TCP stream. In particular, they should judge and maintain a stateful inspection of such packets and flows.

Here are some of the advantages of stateful pattern-matching recognition:

- The capability to directly correlate a specific exploit within a given pattern
- Support for all nonencrypted IP protocols

Systems that perform stateful pattern matching keep track of the arrival order of nonencrypted packets and handle matching patterns across packet boundaries.

However, stateful pattern-matching recognition shares some of the same restrictions as the simple pattern-matching methodology, including an uncertain rate of false positives and the possibility of some false negatives. Additionally, stateful pattern matching consumes more resources in the IPS device because it requires more memory and CPU processing.

Heuristic/Behavior

A different approach to network intrusion detection is to perform heuristic-based analysis. *Heuristic* scanning uses algorithmic logic from statistical analysis of the traffic passing through the network. Its tasks are CPU and resource intensive, so it is an important consideration while planning your deployment. Heuristic-based algorithms may require fine-tuning to adapt to network traffic and minimize the possibility of false positives. For example, a system signature can generate an alarm if a range of ports is scanned on a particular host or network. The signature can also be orchestrated to restrict itself from specific types of packets (for example, TCP SYN packets). Heuristic-based signatures call for more tuning and modification to better respond to their distinctive network environment.

NOTE Heuristic-based intrusion detection rules output alerts in the categories of benign, suspicious, or unknown.

Anomaly

A different practice keeps track of network traffic that diverges from "normal" behavioral patterns. This practice is called *anomaly-based analysis*. The limitation is that what is considered to be normal must be defined. Systems and applications

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

whose behavior can be easily considered as normal could be classified as heuristic-based systems.

However, sometimes it is challenging to classify a specific behavior as normal or abnormal based on different factors, which include the following:

- Negotiated protocols and ports
- Specific application changes
- Changes in the architecture of the network

A variation of this type of analysis is profile-based detection. It allows systems to orchestrate their alarms on alterations in the way that other systems or end users interrelate on the network.

Another kind of anomaly-based detection is protocol-based detection. This scheme is related to, but not to be confused with, the protocol-decode method. The protocol-based detection technique depends on well-defined protocols, as opposed to the protocol-decode method, which classifies any unpredicted value or configuration within a field in the respective protocol as an anomaly. For example, a buffer overflow can be detected when specific strings are identified within the payload of the inspected IP packets.

TIP A buffer overflow occurs when a program attempts to stock more data in a temporary storage area within memory (buffer) than it was designed to hold. This might cause the data to incorrectly overflow into an adjacent area of memory. An attacker could thus craft specific data inserted into the adjacent buffer. Subsequently, when the corrupted data is read, the target computer executes new instructions and malicious commands.

Traditional IDS and IPS provide excellent application layer attack-detection capabilities. However, they do have a weakness. For example, they cannot detect DDoS attacks where the attacker uses valid packets. IDS and IPS devices are optimized for signature-based application layer attack detection. Another weakness is that these systems utilize specific signatures to identify malicious patterns. Yet, if a new threat appears on the network before a signature is created to identify the traffic, it could lead to false negatives. An attack for which there is no signature is called a *zero-day attack*.

Although some IPS devices do offer anomaly-based capabilities, which are required to detect such attacks, they need extensive manual tuning and have a major risk of generating false positives.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

You can use more elaborate anomaly-based detection systems to mitigate DDoS attacks and zero-day outbreaks. Typically, an anomaly detection system monitors network traffic and alerts or reacts to any sudden increase in traffic and any other anomalies. You can also use NetFlow as an anomaly detection tool. NetFlow is a Cisco proprietary protocol that provides detailed reporting and monitoring of IP traffic flows through a network device, such as a router, switch, or the Cisco ASA.

Inline vs. Passive

Intrusion prevention system devices are capable of not only detecting all these security threats but also dropping malicious packets inline. IPS devices may be initially configured in promiscuous mode (passive mode) when you are first deploying them in the network. This is done to analyze the impact to the network infrastructure. Then they are deployed in inline mode to be able to block any malicious traffic in your network.

Figure 19-15 shows how an IPS device is placed inline and drops the noncompliant packet while sending an alert to the monitoring system.

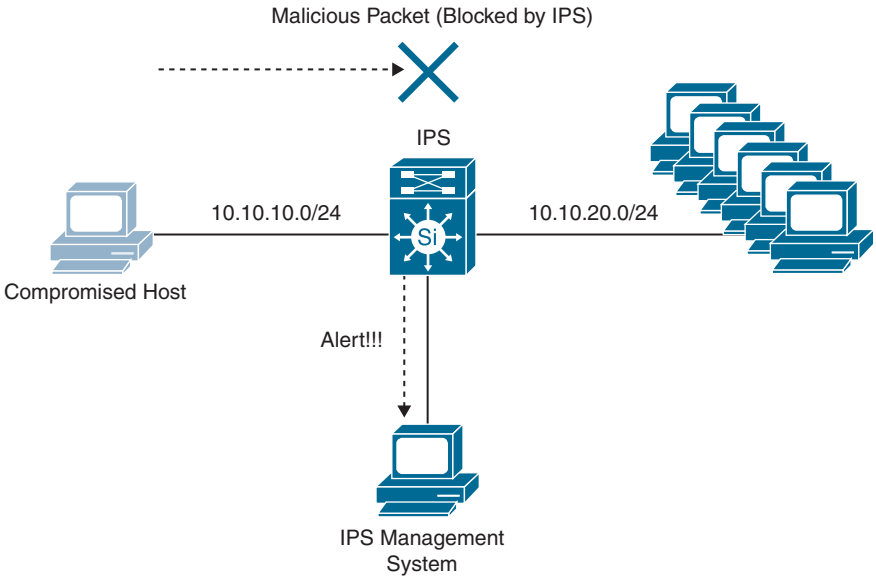


FIGURE 19-15 IPS Device Placed Inline

A few different types of IPSs exist:

- Traditional network-based IPSs (NIPSs)
- Next-generation IPS systems (NGIPSs)
- Host-based IPSs (HIPSs)

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Network-based IDSs and IPSs use several detection methodologies, such as the following:

- Pattern matching and stateful pattern-matching recognition
- Protocol analysis
- Heuristic-based analysis
- Anomaly-based analysis
- Global threat correlation capabilities

HSM

Hardware security modules (HSMs) are physical devices that act as secure cryptoprocessors. This means that they are used for encryption during secure login/authentication processes, during digital signings of data, and for payment security systems. The beauty of a hardware-based encryption device such as an HSM (or a Trusted Platform Module [TPM]) is that it is faster than software encryption.

HSMs can be found in adapter card form, as devices that plug into a computer via USB, and as network-attached devices. They are generally tamper-proof, giving a high level of physical security. They can also be used in high-availability clustered environments because they work independently of other computer systems and are used solely to calculate the data required for encryption keys. However, many of these devices require some kind of management software to be installed on the computer they are connected to. Some manufacturers offer this software as part of the purchase, but others do not, forcing the purchaser to build the management software themselves. Due to this lack of management software, and the cost involved in general, HSMs have seen slower deployment with some organizations. This concept also holds true for hardware-based drive encryption solutions.

Often, HSMs are involved in the generation, storage, and archiving of encrypted key pairs such as the ones used in SSL/TLS sessions online, public key cryptography, and public key infrastructures (PKIs).

Sensors

The sensors in a network are the devices inspecting the traffic, such as an IPS or IDS device. In today's networks, they can also come in the form of a next-generation firewall (NGFW). These devices also have the capability to provide intrusion detection and prevention protections and work as sensors on the network. Additionally, with the use of NetFlow version 9, also known as IPFIX, the network itself can become a sensor.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

In the following sections, we discuss collectors and aggregators, which are used to provide this function by utilizing NetFlow. For instance, it is very difficult to deploy hundreds of sensor appliances in the network of large organizations. In addition, the cost is extremely high. When a security incident or breach is detected, the incident responders need answers fast! They do not have time to go over terabytes of packet captures, and they definitely cannot analyze every computer on the network to find the root cause, miscreant, and source of the breach. You can use NetFlow to obtain a high-level view of what is happening in the network, and then the incident responder can perform a deep-dive investigation with packet captures and other tools later in the investigation. Sniffers then can be deployed as needed in key locations where suspicious activity is suspected. The beauty of NetFlow is that you can deploy it anywhere you have a supported router, switch, firewall, or endpoint.

Collectors

A collector is a physical or virtual appliance that collects network data from infrastructure devices. Collectors can be deployed for many different use cases. On some networks, they are implemented for the purpose of visibility only. For instance, they can be deployed on a network to capture information about the other devices on the network and what version of software they are running. This information is useful to your network teams as well as the InfoSec folks. A collector can also have a specific use case, such as when utilized in NetFlow collection. Figure 19-16 shows a NetFlow collector implemented on a network.

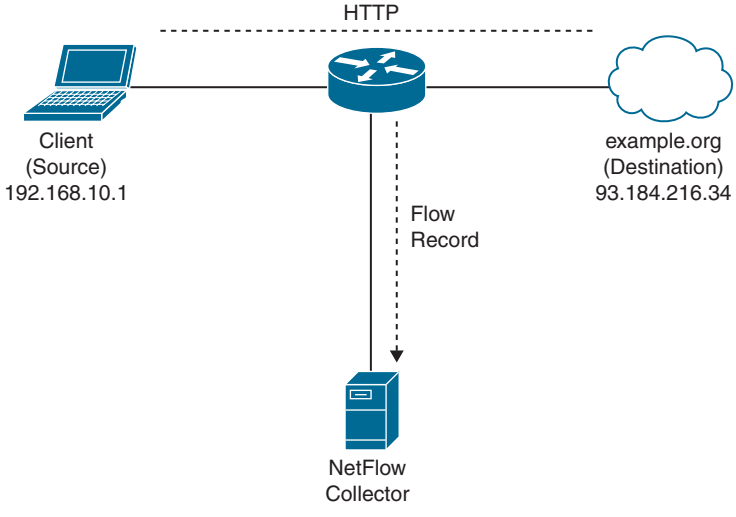


FIGURE 19-16 NetFlow Collector Implementation

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Aggregators

Your security monitoring can be augmented by using a Security Information and Event Management (SIEM) solution. SIEM products combine security event management and security information management. Products such as Splunk and IBM's QRadar offer real-time monitoring of systems and logs, and automation in the form of alerts and triggers. Some of the capabilities of a SIEM solution include data aggregation, which can combine data from network devices, servers, and applications; correlation engines, which automatically look for common attributes of events across the various monitored platforms; compliance with government regulatory auditing processes; and forensic analysis. When you are correlating data, this solution provides for automatic deduplication, or the elimination of redundant data. It may also include scanning for configuration compliance, also known as configuration compliance manager functionality. Utilizing a SIEM solution provides you with features such as dashboards including charts, alerts, graphs, and trends that can be explored deeper by utilizing the drill-down functions.



Firewalls

Nowadays, firewalls are everywhere. Businesses large and small use them, and many households have simpler versions of these protective devices as well. You need to be aware of several types of firewalls, and you definitely want to spend some time configuring hardware and software firewalls. There are many free software-based firewalls and firmware-based emulators that you can download. A quick Internet search provides several options.

The firewall is there to protect the entire network, but other tools are often implemented as well; for example, proxy servers that help protect users and computers by keeping them anonymous; honeypots meant to attract hackers, crackers, and other types of attackers into a false computer or network; and data loss prevention (DLP) devices to keep confidential data from leaving the network. But by far, the most important element in your network is the firewall, so let's begin with that.

Network-based firewalls are primarily used to section off and protect one network from another. They are a primary line of defense and are *extremely* important in network security. There are several types of firewalls; some run as software on server computers, some as standalone dedicated appliances, and some work as just one function of many on a single device. They are commonly represented as a sort of "brick wall" between a LAN and the Internet, as shown in Figure 19-17.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

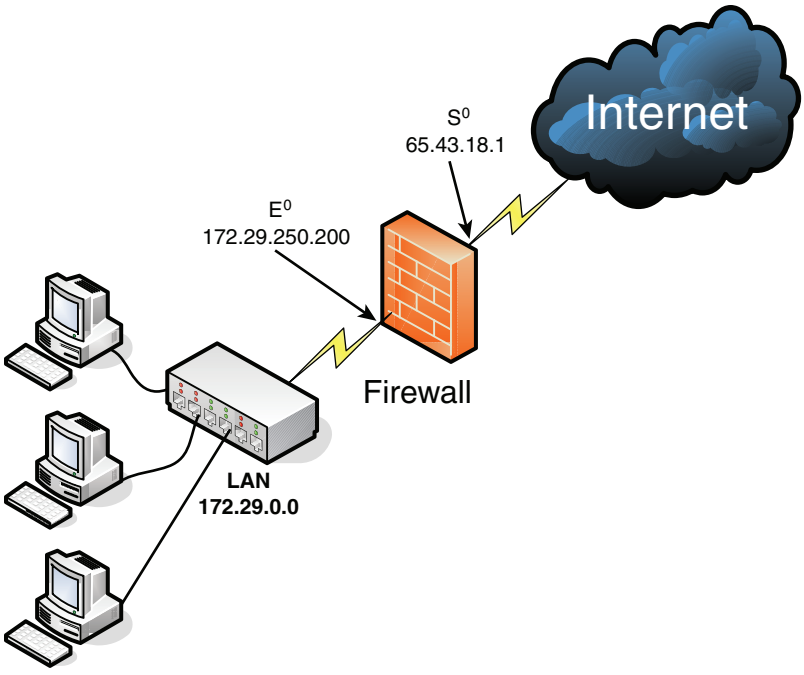


FIGURE 19-17 A Basic Firewall Implementation

Just as a firewall in a physical building is there to slow the spread of a fire and contain it until the fire department arrives, a firewall in a computer network is there to keep fire at bay in the form of malicious attacks. Often, a firewall (or the device the firewall resides on) has NAT in operation as well. In Figure 19-17, note that the firewall has a local address of 172.29.250.200; this connects it to the LAN. It also has an Internet address of 65.43.18.1, enabling connectivity for the entire LAN to the Internet, while hiding the LAN IP addresses. By default, the IP address 65.43.18.1 is completely shielded. This means that all inbound ports are effectively closed and do not enable incoming traffic, unless a LAN computer initiates a session with another system on the Internet. However, a good security administrator always checks this to make sure. First, you should access the firewall's firmware (or software application, as the case may be) and verify that the firewall is on. Next, you scan the firewall with third-party applications such as Nmap (https://nmap.org) or with a web-based port scanning utility. If any ports are open or unshielded, you should deal with them immediately. Then you should rescan the firewall for vulnerabilities.

Important point: Firewalls should be used only as they were intended. The company firewall should not handle any other extraneous services—for example, acting as a web server or SMTP server. When a firewall is used as it was intended, its vulnerability is reduced.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Generally, a firewall inspects traffic that passes through it and permits or denies that traffic based on rules set by an administrator. These rules are stored within access control lists (ACLs). With regard to firewalls, an ACL is a set of rules that apply to a list of network names, IP addresses, and port numbers. These rules can be configured to control inbound and outbound traffic. This is a bit different from ACLs with respect to operating systems, but the same basic principles apply: Basically, one entity is granted or denied permission to another entity. If you decide that a specific type of traffic should be granted access to your network, you would *explicitly allow* that traffic as a rule within an ACL. If, on the other hand, you decide that a specific type of traffic should *not* be granted access, you would *explicitly deny* that traffic within an ACL. And finally, if a type of network traffic is not defined in the firewall's rule set, it should be stopped by default. This is the concept of *implicit deny* and is usually a default rule found in a firewall's ACL. It is often added automatically to the end of a firewall's rule set (ACLs) and is also known as "block all."

Firewall rules should be specific. Here's an example of a firewall rule:

```
deny TCP any any port 53
```

This rule can be used to restrict DNS zone transfers (as they run on top of TCP and use port 53), but other DNS traffic still functions properly. The rule is specific; it gives the transport layer protocol to be filtered and the exact port, and also states that it applies to *any* computer's IP address on the inbound and outbound side. Be careful with firewall rules and ACLs; they need to be written very cautiously so as not to filter required traffic.

NOTE Traffic can also be passed to other computers and servers, or to specific ports.

A lot of today's firewalls have two types of firewall technologies built into them: SPI and NAT. However, you also should be aware of a few other types of firewall methodologies:

- **Packet filtering:** Inspects each packet passing through the firewall and accepts or rejects it based on rules. However, there are two types: *stateless packet inspection* and *stateful packet inspection* (also known as SPI or a stateful firewall). A *stateless* packet filter, also known as pure packet filtering, does not retain memory of packets that have passed through the firewall; due to this, a *stateless* packet filter can be vulnerable to IP spoofing attacks. But a firewall running stateful packet inspection is normally not vulnerable to this because it keeps track of the state of network connections by examining the header in each packet. It can distinguish between legitimate and illegitimate packets. This function operates at the network layer of the OSI model.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

- **NAT gateway:** Also known as NAT endpoint filtering, filters traffic according to ports (TCP or UDP). This can be done in three ways: by way of basic endpoint connections, by matching incoming traffic to the corresponding outbound IP address connection, or by matching incoming traffic to the corresponding IP address and port.
- **Application-level gateway (ALG):** Applies security mechanisms to specific applications, such as FTP or BitTorrent. It supports address and port translation and checks whether the type of application traffic is allowed. For example, your company might allow FTP traffic through the firewall but might decide to disable Telnet traffic (probably a wise choice). The ALG checks each type of packet coming in and discards Telnet packets. Although this adds a powerful layer of security, the price is that it is resource-intensive, which could lead to performance degradation.
- **Circuit-level gateway:** Works at the session layer of the OSI model and applies security mechanisms when a TCP or UDP connection is established; it acts as a go-between for the transport and application layers in TCP/IP. After the connection has been made, packets can flow between the hosts without further checking. Circuit-level gateways hide information about the private network, but they do not filter individual packets.

A firewall can be set up in several different physical configurations. For example, it could be set up in a back-to-back configuration (two firewalls surrounding the screened subnet or DMZ), as shown in Figure 19-18, or as a three-leg perimeter configuration.

Key Topic

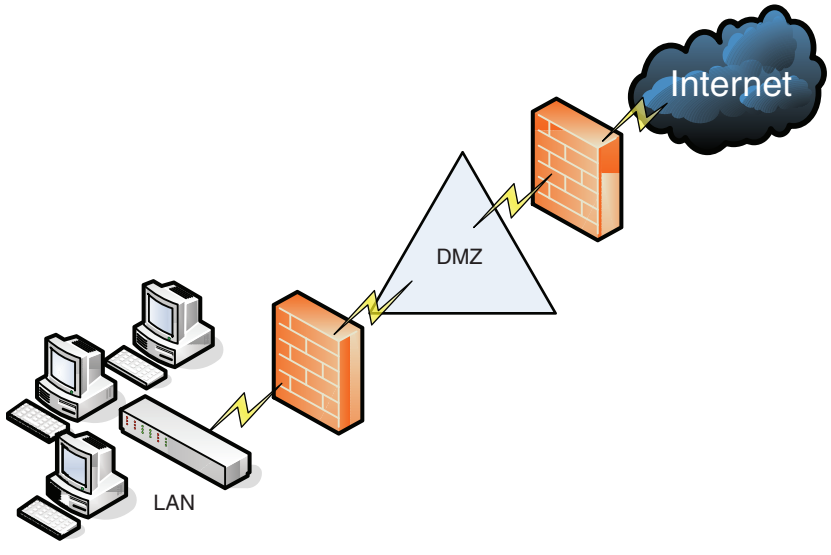


FIGURE 19-18 Back-to-Back Firewall/DMZ (Screened Subnet) Configuration

Generally, there is one firewall with the network and all devices and computers residing "behind" it. By the way, if a device is "behind" the firewall, it is also considered to be "after" the firewall, and if the device is "in front of" the firewall, it is also known as being "before" the firewall. Think of the firewall as the drawbridge of a castle. When you are trying to gain admittance to the castle, the drawbridge is probably closed. You would be in front of the drawbridge, and the people inside the castle would be behind the drawbridge. This is a basic analogy but should help you understand the whole "in front of" and "behind" business as it relates to data attempting to enter the network and devices that reside on your network.

Logging is also important when it comes to a firewall. Firewall logs should be the first thing you check when an intrusion has been detected. You should know how to access the logs and how to read them. For example, Figure 19-19 shows two screen captures: The first displays the Internet sessions on a basic small office/home office (SOHO) router/firewall, and the second shows log events such as blocked packets. Look at the blocked Gnutella packet that is pointed out. You know it is a Gnutella packet because the inbound port on the firewall that the external computer is trying to connect to shows as port 6346; this associates with Gnutella. Gnutella is an older P2P file-sharing network. None of the computers on this particular network use or are in any way connected to the Gnutella service. These external computers are just random clients of the Gnutella P2P network trying to connect to anyone possible.

It's good that these packets have been blocked, but maybe you don't want the IP address shown (24.253.3.20) to have any capability to connect to your network at all. To eliminate that IP, you could add it to an inbound filter or to an ACL.

So far, we have discussed host-based and network-based firewalls. However, both of these types of firewalls can also fall into the category of application firewall. If either type runs protocols that operate on the application layer of the OSI model, it can be classified as an application firewall. That means that it can control the traffic associated with specific applications. This is something a stateful network firewall cannot do because this function operates at the application layer of the OSI model. Many host-based firewalls fall into this category, but when it comes to network-based firewalls, it varies. A basic SOHO router with built-in firewalling capabilities would usually not fall into the application firewall category. However, more advanced network appliances from companies such as Barracuda, Citrix, Fortinet, and Smoothwall do fall into this category. This means that they allow for more in-depth monitoring of the network by controlling the input, output, and access to applications and services all the way up through the application layer of the OSI model. These appliances might also be referred to as *network-based application layer firewalls*. Now that's a mouthful—just be ready for multiple terms used by companies and technicians.

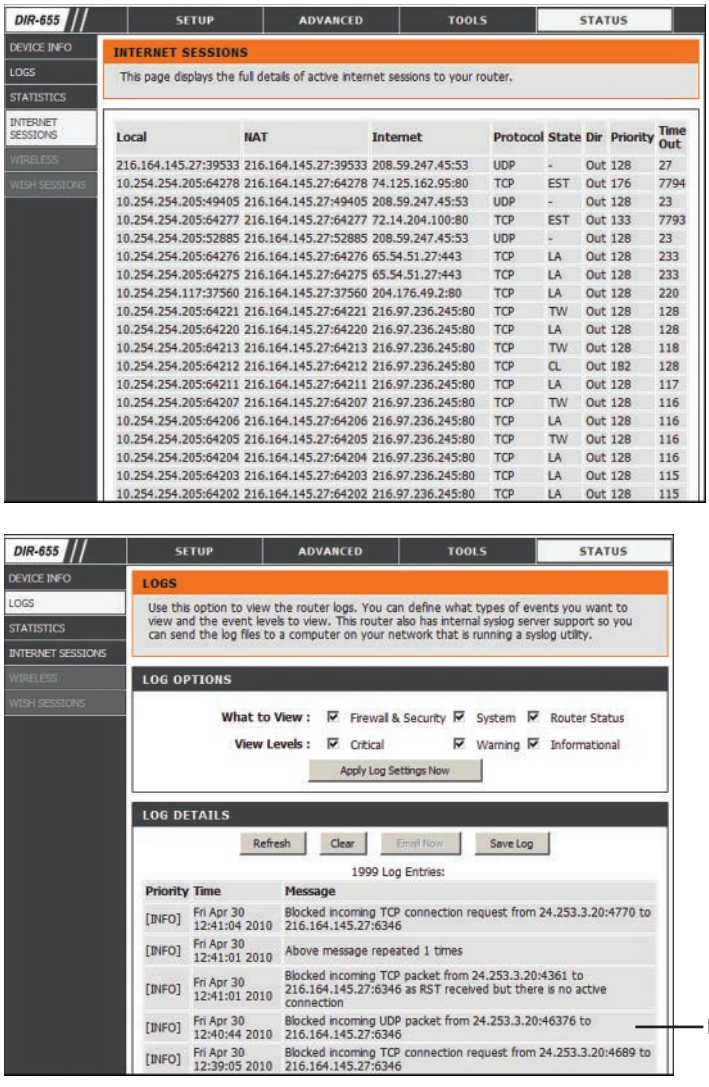


FIGURE 19-19 SOHO Router/Firewall Internet Sessions

Going a step further, some of the aforementioned network appliances have tools that are designed to specifically protect HTTP sessions from XSS attacks and SQL injection. These types of tools are known as ***web application firewalls***. WAFs can help protect the servers in your environment.

NOTE A firewall appliance needs more than one network adapter so that it can connect to more than one network; this is known as a *multihomed connection*. It might be dual-homed (two adapters), or perhaps it has more, maybe three network adapters, in case you want to implement a DMZ (screened subnet) or another perimeter security technique.

Firewalls are often considered to be all-in-one devices, but actually they provide specific functionality as discussed in this section. Still, it is common to hear people refer to a firewall when they are really talking about another technology or even another device. For example, many SOHO users have an all-in-one multifunction network device. This device has four ports for wired connections, plus a wireless antenna; it connects all the computers to the Internet and also has a firewall built in. Because some users consider this to be simply a firewall, you should teach them about the benefits of disabling service set identifier (SSID) broadcasting and enabling MAC filtering. By disabling SSID broadcasting, the average user cannot connect wirelessly to the device. An attacker knows how to bypass this setting, but it is an important element of security that you should implement after all trusted computers have been connected wirelessly. MAC filtering denies access to any computer that does not have one of the MAC addresses you list.

To make matters a bit more confusing, a firewall can also act as, or in combination with, a proxy server, which we discuss in the following section.

A relatively newer concept, ***unified threat management (UTM)*** is the culmination of everything we have discussed in this chapter so far. As early as the year 2000, it was realized that the firewall was no longer enough to protect an organization's network. Other devices and technologies such as NIDS/NIPS, content filters, antimalware gateways, data leak prevention, and virtual private networks were added to the network to better protect it. However, with all these extra devices and technologies come added cost and more administration. So, UTM providers simplify the whole situation by offering all-in-one devices that combine the various levels of defense into one solution. The all-in-one device might also be referred to as a ***next-generation firewall (NGFW)***. Companies such as Cisco, Fortinet, and Sophos (to name a few) offer UTM and NGFW solutions; often this is a single device that sits last on the network before the Internet connection. They usually come with a straightforward web-based GUI, which is good news for the beleaguered security administrator like you who might be burning the midnight oil researching the latest attacks and prevention methods.

There's a caveat to all this, and it is a common theme in network security: a single point of defense is a single point of failure. Get past the UTM, and the attacker's job is done. Secondary and backup UTM devices, as well as server-based HIDSs, strike

a balance and create a certain level of defense in depth, while still retaining a level of simplicity. Another consideration is that UTMs should be quick. If they are to take the place of several other devices, their data processing and traffic flow requirements are steep. As a smart network or security administrator, you will consider a device that exceeds your current needs and then some.

Another example of a proxy in action is Internet ***content URL/filtering***. An Internet content/URL filter, or simply a content filter, is usually applied as software at the application layer and can filter out various types of Internet activities such as websites accessed, email, instant messaging, and more. It often functions as a content inspection device, and disallows access to inappropriate web material (estimated to be a big percentage of the Internet!) or websites that take up far too much of an organization's Internet bandwidth. Internet content filters can be installed on individual clients, but by far the more efficient implementation is as an individual proxy that acts as a mediator between all the clients and the Internet. These proxy versions of content filters secure the network in two ways:

- By forbidding access to potentially malicious websites
- By blocking access to objectionable material that employees might feel is offensive

They can also act as URL filters; even if employees inadvertently type an incorrect URL, they can rest assured that any objectionable material will not show up on their display.

Internet filtering appliances analyze just about all the data that comes through them, including Internet content, URLs, HTML tags, metadata, and security certificates such as the kind you would automatically receive when going to a secure site that starts with *https*. However, revoked certificates and certificate revocation lists (CRLs) are not filtered because they are only published periodically. Some of these appliances are even capable of malware inspection.

Another similar appliance is the web security gateway. Web security gateways (such as Forcepoint, previously known as Websense) act as go-between devices that scan for viruses, filter content, and act as data loss prevention (DLP) devices. This type of content inspection/content filtering is accomplished by actively monitoring the users' data streams in search of malicious code, bad behavior, or confidential data that should not be leaked outside the network.

Next, we compare hardware versus software firewalls and appliance versus host-based versus virtual firewalls.

Hardware vs. Software

When it comes to firewall devices, they can come in many forms and flavors and include a variety of different features. Most organizations have at least one hardware firewall, the traditional firewall. This purpose-built device runs proprietary software to provide the firewall functions discussed in previous sections. Hardware-based firewalls are normally sold and differentiated based on the throughput capacity as well as the number and type of physical ports available. For instance, a lower-end firewall might support only 100 Mbps of throughput; however, a higher-end service-provider-grade hardware firewall might support 300 Gbps throughput. This is all based on the memory and CPU as well as other proprietary hardware that makes up the device. On the other hand, a software-based firewall would be a piece of software that can run on a specific platform to provide the firewall functions discussed previously. For instance, a software firewall may be able to run on a supported Linux distribution, utilizing whatever hardware that platform is installed on. Typically these days, a hardware-based firewall is best for higher performance.

Appliance vs. Host-based vs. Virtual

An appliance is essentially the same thing as a hardware-based firewall. It again is purpose built to run a specific piece of software or firmware that provides the firewall functions. All of the same assertions mentioned for software-based firewalls would apply to an appliance-based firewall as well.

Host-based firewalls are often referred to as "personal firewalls." Personal firewalls are software applications that you can install on end-user machines or servers to protect them from external security threats and intrusions. The term *personal firewall* typically applies to basic software that can control Layer 3 and Layer 4 access to client machines. Most modern operating systems have host-based firewalls built in and enabled by default. It is important to implement an on-by-default firewall for your user devices as well as servers.

A virtual firewall is typically used in a virtual environment for the purpose of controlling access between virtual machines. It can also come in the form of a container that would enable you to integrate into a containerized environment and control traffic within. Figure 19-20 illustrates a virtual firewall implementation.

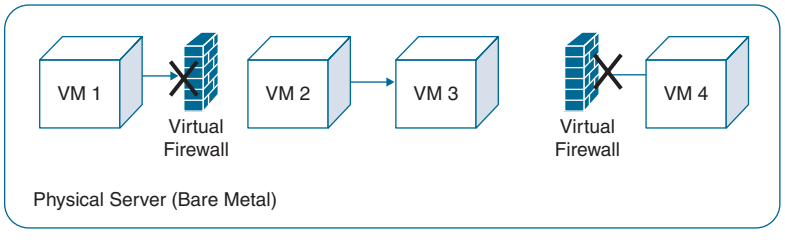


FIGURE 19-20 A Virtual Firewall Implementation

We discussed each of the tools and technologies separately in this chapter so that you understand how to work with each. But you should keep in mind that many of these technologies are consolidated into a single solution, a trend that will likely continue as we move forward.

Key Topic Access Control List

The most basic implementation of an access control is an access control list (ACL). When an ACL is applied to network traffic, it is called a *network ACL*. Cisco networking devices such as routers, switches, and firewalls include network ACL capabilities to control access to network resources. As for port-based access controls, network ACLs and firewalling are usually seen as special cases of the ABAC model and also sometimes classified as identity-based or rule-based access control because they base the control decision on attributes such as IP or MAC addresses or Layer 4 information. Security group ACLs, on the other hand, are access lists based on the role of the subject trying to access a resource, and they implement role-based access control.

Network ACLs can be implemented at various levels of the OSI model:

- A Layer 2 ACL operates at the data link layer and implements filters based on Layer 2 information. An example of this type of access list is a MAC access list, which uses information about the MAC address to create the filter.
- A Layer 3 ACL operates at the networking layer. Cisco devices usually allow Layer 3 ACLs for different Layer 3 protocols, including the most used ones nowadays: IPv4 and IPv6. In addition to selecting the Layer 3 protocol, a Layer 3 ACL allows the configuration of filtering for a protocol using raw IP, such as OSPF or ESP.
- A Layer 4 ACL operates at the transport layer. An example of a Layer 4 ACL is a TCP- or UDP-based ACL. Typically, a Layer 4 ACL includes the source and destination. This allows filtering of specific upper-layer packets.

Key Topic Route Security

There are different route manipulation attacks, but one of the most common is the BGP hijacking attack. BGP is a dynamic routing protocol used to route Internet traffic. An attacker can launch a BGP hijacking attack by configuring or compromising an edge router to announce prefixes that have not been assigned to his or her organization. If the malicious announcement contains a route that is more specific

than the legitimate advertisement or presents a shorter path, the victim's traffic may be redirected to the attacker. In the past, threat actors have leveraged unused prefixes for BGP hijacking to avoid attention from the legitimate user or organization. Additionally, an on-path attack can occur at Layer 3 by placing a rogue router on the network and then tricking the other routers into believing that this new router has a better path. This could cause network traffic to flow through the rogue router and again allow the attacker to steal network data. You can mitigate attacks such as these in various ways, including using routing authentication protocols and filtering information from being advertised or learned on specific interfaces.

Quality of Service

Key Topic

Administrators can use quality of service (QoS) capabilities to control application prioritization. Protocol discovery features show you the mix of applications currently running on the network. This helps you define QoS classes and policies, such as how much bandwidth to provide to mission-critical applications and how to determine which protocols should be policed. Per-protocol bidirectional statistics are available, such as packet and byte counts, as well as bit rates.

After you classify the network traffic, you can apply the following QoS features:

- Apply class-based weighted fair queuing (CBWFQ) for guaranteed bandwidth
- Enforce bandwidth limits using policing
- Mark for differentiated service downstream or from the service provider using the type of service (ToS) bits or DSCPs in the IP header
- Drop policy to avoid congestion using weighted random early detection (WRED)

Implications of IPv6

Key Topic

Most internal networks (meaning LANs) are either subnetted or are classless in nature. It's important to know the private IP ranges and their classes, but just remember that classless is very common, especially in larger networks.

You should also know the categories of IPv6 addresses. Table 19-4 provides a review of these types. Keep in mind that the standard "private" range for IPv6 is FE80::/10, which spans addresses that start with FE80, FE90, FEA0, and FEB0. This is the default reserved range of IPv6 addresses that computers on a LAN (and behind a firewall) are assigned from.

Table 19-4 Types of IPv6 Addresses

IPv6 Type	Address Range	Purpose
Unicast	Global unicast starts at 2000 Link-local ::1 and FE80::/10	Address assigned to one interface of one host.
Anycast	Structured like unicast addresses	Address assigned to a group of interfaces on multiple nodes. Packets are delivered to the "first" interface only.
Multicast	FF00::/8	Address assigned to a group of interfaces on multiple nodes. Packets are delivered to all interfaces.

There are risks involved with IPv6 autoconfigured addresses. When a computer is physically connected to a network, it can easily obtain an IPv6 address without any user interaction and begin communicating with other hosts on the network, perhaps in an insecure manner. To avoid this situation, you should consider using 802.1X authentication, which stops IPv6 traffic until the host is authenticated to the network. You can read more on 802.1X in Chapter 20, "Installing and Configuring Wireless Security Settings," and Chapter 24, "Implementing Authentication and Authorization Solutions." Also, consider using encrypted tunnels and network adapters that are certified for secure wired and/or wireless transmissions.

A last word about IP security: Both IPv4 and IPv6 have security issues, and both have various ways that they can be secured. Don't be fooled; both types of IP networks need to be designed with security in mind. For example, IPv6 has IPsec support built in, but that might not be the best method of security for your organization. IPv4 also can make use of IPsec, and in that respect can be just as secure as IPv6, but the support isn't built in, so you might choose to implement alternative security methods for IPv4. Many networks use both protocols, and though one is working in a secure manner, that doesn't mean the other protocol is protected. Remember to design both types of IP networks to address all security concerns, and test them thoroughly on multiple platforms.



Port Spanning/Port Mirroring

To capture packets and measure throughput, you need a tap on the network before you can start monitoring. Most tools that collect throughput leverage a single point configured to provide raw data, such as pulling traffic from a switch or router. If the access point for the traffic is a switch, typically a network port is configured as a Switched Port Analyzer (SPAN) port, sometimes also called port mirroring or port

spanning. The probe capturing data from a SPAN port can be either a local probe or data from a SPAN port that is routed to a remote monitoring tool.

Example 19-2 demonstrates how to configure a Cisco switch port as a SPAN port so that a collection tool can be set up to monitor throughput. The SPAN session for this example is ID 1 and is set to listen on the fastEthernet0/1 interface of the switch while exporting results to the fastEthernet0/10 interface of the switch. SPAN sessions can also be configured to capture more than one VLAN.

Example 19-2 Configuring a SPAN Port

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface fastEthernet0/1
Switch(config)# monitor session 1 destination interface fastEther-
net0/10 encapsulation dot1q
Switch(config)# end
```

Another method for capturing packets is to place a ***port tap*** in the line of traffic. This is a common tactic for monitoring throughput from a routing point or security checkpoint, such as a firewall configured to have all traffic cross it for security purposes.

Packet capture tools, sometimes referred to as packet analyzers or packet sniffers, are hardware or programs that intercept and log traffic as it passes over a digital network. This happens as a data stream crosses the network while the packet capture tool captures each packet and decodes the packet's raw data elements representing the various fields that can be interpreted to understand its content.

The requirements for capturing packets are having access to the data, the capability to capture and store the data, tools to interoperate with the data, and capabilities to use the data. This means the tool must be able to access the data from a port tap and have enough storage to collect the data, and it must be able to read the results and have a method to use those results for some goal. Failing at any of these will most likely result in an unsuccessful packet capture session.



Monitoring Services

To operate a clean, secure network, you must keep an eye on your systems, applications, servers, network devices, power devices, and the entire network in general. One way to do this is to monitor the network. This surveillance of the network in and of itself increases the security of your entire infrastructure. By periodically watching everything that occurs on the network, you become more familiar with day-to-day happenings and over time get quicker at analyzing whether an event is

legitimate. It can help to think of yourself as Hercule Poirot, the Belgian detective—*seeing* everything that happens on your network, and ultimately *knowing* everything that happens. It might be a bit egotistical sounding, but whoever said that IT people don't have an ego?

This surveillance can be done in one of two ways: manual monitoring or automated monitoring. When manually monitoring the network, you are systematically viewing log files, policies, permissions, and so on. But this process can also be automated. For example, several available data mining programs can automatically sift through logs and other files for the exact information you want to know. In addition, applications such as antivirus, intrusion detection systems (IDSs), and intrusion prevention systems (IPSs) can automatically scan for errors, malicious attacks, and anomalies. The three main types of automated monitoring are signature-based, anomaly-based, and behavior-based.

All the methodologies in the world won't help you unless you know how to use some monitoring tools and how to create baselines. By using performance monitoring gizmos and software, incorporating protocol analyzers, and using other analytical utilities in the GUI and the command line, you can really "watch" the network and quickly mitigate threats as they present themselves.

In the following sections, we use the Performance tool in Windows, the Wireshark protocol analyzer, and other analytical tools within the command line and the GUI. These are just a few examples of performance and network monitoring tools out there, but they are commonly used in the field and should give you a decent idea of how to work with any tools in those categories.

Performance Baselining

Baselining is the process of measuring changes in networking, hardware, software, applications, and so on. Documenting and accounting for changes in a baseline is known as ***baseline reporting***. Baseline reporting enables you, as security administrator, to identify the ***security posture*** of an application, system, or network. The security posture can be defined as the risk level to which a system, or other technology element, is exposed. ***Security posture assessments (SPAs)*** use baseline reporting and other analyses to discover vulnerabilities and weaknesses in systems.

NOTE Ultimately, a security baseline defines the basic set of security objectives that must be met by any given service or system. Baseline deviation, meaning any changes or discrepancies as compared to the baseline, should be investigated right away, especially for mission-critical systems and RTOS-based computers.

Let's get into baselining a little more and show one of the software tools you can use to create a baseline.

Creating a baseline consists of selecting something to measure and measuring it consistently for a period of time. For example, you might want to know what the average hourly data transfer is to and from a server's network interface. There are a lot of ways to measure this, but you could possibly use a performance monitoring tool or a protocol analyzer to find out how many packets cross through the server's network adapter. This could be run for one hour (during business hours, of course) every day for two weeks. Selecting different hours for each day would add more randomness to the final results. By averaging the results together, you get a baseline. Then you can compare future measurements of the server to the baseline. This helps you define what the standard load of your server is and the requirements your server needs on a consistent basis. It also helps when installing other like computers on the network. The term *baselining* is most often used to refer to monitoring network performance, but it actually can be used to describe just about any type of performance monitoring and benchmarking. The term *standard load* is often used when referring to servers. A configuration baseline defines what the standard load of the server is for any measured objects. When it comes to performance-monitoring applications, objects are all of the components in the server (for example, CPU, RAM, hard disk, and so on). They are measured using counters. A typical counter would be the % Processor Time of the CPU. This is used by the Task Manager in Windows systems.

An example of one of these tools is the Performance Monitor tool in Windows (see Figure 19-21). It can help to create baselines measuring network activity, CPU usage, memory, hard drive resources used, and so on. It should also be used when monitoring changes to the baseline. The program works basically the same in all versions of Windows, be it client or server, but the navigation to the program varies. To simplify matters, you can go to the Run prompt and type **perfmon.exe** in Windows to open the program.

The CPU is probably the most important component of the computer. In Figure 19-21, the CPU counter hit 100 percent once, which is normal when opening applications or starting other processes. However, if the CPU maxes out often, a percentage of clients cannot obtain access to resources on the computer. If the computer is a server, that means trouble. CPU spiking could be due to normal usage, or it could be due to malicious activity or perhaps bad design. Further analysis would be necessary to determine the exact cause. If the system is a virtual machine, there is a higher probability of CPU spikes. Proper design of VMs is critical, and they must have a strong platform to run on if they are to serve clients properly. Known as a counter, the CPU % Processor Time is just one of many counters. A smart security auditor measures the activity of other objects such as the hard drive, paging file, memory (RAM), network adapter, and whatever else is specific to the organization's needs. Each object has several counters to select from. For example, if you are analyzing a

web server, you would probably want to include the HTTP Service Request Queries object, and specifically the ArrivalRate and CurrentQueueSize counters in your examination.

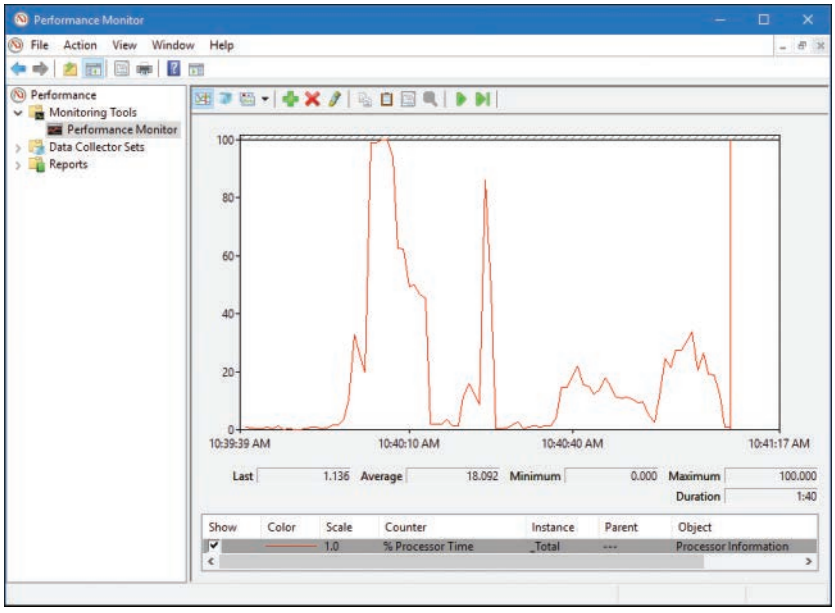


FIGURE 19-21 Performance Monitor in Windows

Figure 19-21 shows the Performance Monitor screen, but this only gives you a brief look at the system. The window of time is only a minute or so before the information refreshes. However, you can record this information over *x* periods of time and create reports from the recorded information. By comparing the Performance Monitor reports and logs, you ultimately create the baseline. The key is to measure the same way at the same time each day or each week. This provides accurate comparisons. However, you should keep in mind that performance recording can be a strain on resources. You should verify that the computer in question can handle the tests first before you initiate them.

Making reports is all fine and good (and necessary), but it is wise to also set up alerts. Alerts can be generated automatically by the system and sent to administrators and other important IT people. These alerts can be set off in a myriad of ways, all of your choosing—for example, if the CPU were to trip a certain threshold or run at 90 percent for more than a minute (although this is normal in some environments). Or maybe the physical disk was peaking at 100 Mbps for more than five minutes. If these types of things happen often, the system should be checked for malicious activity, illegitimate usage, or the need for an upgrade.

A tool similar to Performance Monitor used in Linux systems is called System Monitor. The different versions of Linux also have many third-party tools that you can use for performance monitoring. macOS uses Activity Monitor.

File Integrity Monitors

File integrity is also important when securing log files. Encrypting the log files through the concept known as hashing is a good way to verify the integrity of the log files if they are moved and/or copied. Finally, you could flat-out encrypt the entire contents of the file so that other users cannot view it. Integrity means that data has not been tampered with. Authorization is necessary before data can be modified in any way; this is done to protect the data's integrity. For example, if a person deletes a required file, either maliciously or inadvertently, the integrity of that file is violated. There should be permissions in place to stop the person from deleting the file.

Here's a tip for you: Some organizations do not delete data—ever! Loaded on an individual computer, a host-based intrusion detection system (HIDS) analyzes and monitors what happens inside that computer—for example, whether any changes have been made to file integrity. File integrity monitoring (FIM) should be implemented on all mission-critical systems. One very well known FIM solution is Tripwire.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 19-5 lists a reference of these key topics and the page number on which each is found.

Table 19-5 Key Topics for Chapter 19

Key Topic Element	Description	Page Number
Section	Load Balancing	488
Section	Network Segmentation	489
Figure 19-1	Firewall providing network segmentation	489
Figure 19-4	Screened Subnet	491

Key Topic Element	Description	Page Number
Section	Virtual Private Network	494
Section	Remote Access vs. Site-to-Site	496
Section	DNS	509
Section	Network Access Control	510
Section	Out-of-Band Management	510
Section	Port Security	511
Section	Bridge Protocol Data Unit Guard	512
Section	Dynamic Host Configuration Protocol Snooping	512
Section	Network Appliances	513
Section	Firewalls	526
Figure 19-18	Back-to-Back Firewall/Screened Subnet Configuration	529
Section	Access Control List (ACL)	535
Section	Route Security	535
Section	Quality of Service	536
Section	Implications of IPv6	536
Section	Port Spanning/Port Mirroring	537
Section	Monitoring Services	538
Section	File Integrity Monitors	542

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

virtual private network (VPN), VPN concentrator, always-on VPN, split tunneling, full tunnel, site-to-site VPNs, remote-access VPNs, DHCP snooping, forward proxy, reverse proxy, network-based intrusion detection system (NIDS), promiscuous mode, network-based intrusion prevention system (NIPS), false positive, false negative, web application firewalls, unified threat management (UTM), next-generation firewall (NGFW), content URL/filtering, port tap, baselining, baseline reporting, security posture, security posture assessments (SPAs)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. During an audit of your servers, you notice that most servers have large amounts of free disk space and have low memory utilization. What is the primary impact on your organization when utilizing this type of practice?
2. What concept of network security can help defend against pivoting during a compromise?
3. Which kind of VPN implementation does not install software on the host system to establish the VPN connection?
4. In which type of attack does an attacker generate a high number of requests over port 53?
5. Which type of network security control could be used to control access to a network based on the security posture?
6. Which type of management access would require an alternative path?
7. Which feature present in most Cisco switches is used to provide access control by restricting the MAC address that can connect?
8. Which type of appliance secures a network by keeping machines behind it anonymous?
9. Which type of ACL would control access based on MAC address?
10. Which routing protocol is most common in route manipulation attacks?
11. Which QOS feature can be used for guaranteed bandwidth?
12. Which type of IPv6 address is structured like a unicast address?

This page intentionally left blank



This chapter covers the following topics related to Objective 3.4 (Given a scenario, install and configure wireless security settings) of the CompTIA Security+ SY0-601 certification exam:

- Cryptographic protocols
  - WiFi Protected Access 2 (WPA2)
  - WiFi Protected Access 3 (WPA3)
  - Counter-mode/CBC-MAC protocol (CCMP)
  - Simultaneous Authentication of Equals (SAE)
- Authentication protocols
  - Extensible Authentication Protocol (EAP)
  - Protected Extensible Application Protocol (PEAP)
  - EAP-FAST
  - EAP-TLS
  - EAP-TTLS
  - IEEE 802.1X
  - Remote Authentication Dial-in User Service (RADIUS) Federation
- Methods
  - Pre-shared key (PSK) vs. Enterprise vs. Open
  - WiFi Protected Setup (WPS)
  - Captive portals
- Installation considerations
  - Site surveys
  - Heat maps
  - WiFi analyzers
  - Channel overlaps
  - Wireless access point (WAP) placement
  - Controller and access point security

Installing and Configuring Wireless Security Settings

This chapter briefly digs into the topic of installing and configuring wireless security settings. We start out by discussing cryptographic protocols such as WPA2, WPA3, and CCMP, as well as Simultaneous Authentication of Equals (SAE). From there we cover authentication protocols including Extensible Authentication Protocol (EAP), Protected Extensible Authentication Protocol (PEAP), EAP-FAST, EAP-TLS, EAP-TTLS, and IEEE 802.1X, as well as a quick discussion on Remote Authentication Dial-in User Service (RADIUS) Federation. You also learn about wireless security methods such as preshared key (PSK) versus Enterprise and Open, as well as an overview of Wi-Fi Protected Setup (WPS) and captive portals. Finally, you learn about wireless installation considerations. This section includes a discussion about site surveys, heat maps, Wi-Fi analyzers, channel overlaps, wireless access point (WAP) placement, and controller access point security.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 20-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 20-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Cryptographic Protocols	1–2
Authentication Protocols	3–7
Methods	8–9
Installation Considerations	10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What WPA version utilizes Simultaneous Authentication of Equals (SAE)?
   a. WPA3
   b. WPA2
   c. WPA
   d. None of these answers are correct.

2. Which protocol is based on AES and provides stronger security over TKIP?
   a. WPA
   b. CCMP
   c. WPA2
   d. All of these answers are correct.

3. Which of the following is a data link layer authentication technology that defines port-based access control?
   a. 802.1X
   b. 802.11ac
   c. 802.11a
   d. None of these answers are correct.

4. Which of the following components of 802.1X would be an access point or switch?
   a. Authorization server
   b. Authentication server
   c. Supplicant
   d. Authenticator

5. Which of the following components of 802.1X runs on a client workstation?
   a. Authorization server
   b. Authentication server
   c. Authenticator
   d. Supplicant

6. Which of the following components of 802.1X includes an authentication database?
   a. EAP-FAST
   b. EAP-TTLS
   c. EAP-TLS
   d. PEAP

7. Which type of EAP authentication uses Secure Tunneling?
   a. EAP-TLS
   b. EAP-FAST
   c. EAP-TTLS
   d. PEAP

8. _____is automatically used when you select WPA-Personal.
   a. PKI
   b. WPA
   c. PSK
   d. Open

9. Which of the following was created to give users easy connectivity to wireless access points?
   a. WPA
   b. WPS
   c. WEP
   d. All of these answers are correct.

10. Which of the following can help with the proper placement of access points?
    a. WEP
    b. Channel overlaps
    c. RADIUS
    d. Site survey

Foundation Topics

Cryptographic Protocols

Key Topic

Weak encryption or no encryption can be the worst thing that can happen to a wireless network. This lack of encryption can occur for several reasons—for example, if someone wanted to connect an older device or a device that hasn't been updated, and that device can run only a weaker, older type of encryption. It's important to have strong encryption in your network; as of this writing, ***Wi-Fi Protected Access 2 (WPA2)*** is the most widely used wireless protocol. It can be used with TKIP or, better yet, AES. Remember that the encryption level of the WAP and the encryption level of the network adapters that connect to it need to be the same. ***Wi-Fi Protected Access 3 (WPA3)*** is the future of Wi-Fi security and is quickly becoming the replacement for WPA2. WPA3 includes a more robust authentication mechanism over WPA2. It also provides a higher level of encryption capabilities. It enables a very robust authentication based on passwords by utilizing a technology called ***Simultaneous Authentication of Equals (SAE)***. This innovation in Wi-Fi security replaces the preshared key (PSK). SAE helps protect against brute-force password attacks and offline dictionary attacks. The sections that follow describe these cryptographic protocols in more detail.

Wi-Fi Protected Access 2 (WPA2)

Wi-Fi Protected Access 2, or WPA2 as it is commonly referred to, is the next evolution of Wi-Fi Protected Access (WPA) that was created to enhance and replace the original WPA protocol. All three current WPA versions were developed by the Wi-Fi Alliance, which is an organization consisting of a partnership between many companies in the industry with the goal of standardizing on Wi-Fi technologies. WPA2 includes support for Counter-mode/CVC-MAC protocol (CCMP) and AES-based encryption. WPA2 can be used in PSK or Enterprise mode. PSK is typically used in personal home deployments of Wi-Fi networks because it is much easier to implement and requires only a single key or passphrase to be configured on the access point and client for authentication. Enterprise mode requires the use of a back-end authentication server such as a RADIUS server.

Wi-Fi Protected Access 3 (WPA3)

As stated previously, Wi-Fi Protected Access 3 (WPA3) was developed as the eventual replacement of WPA2. It was also developed by the Wi-Fi Alliance, which originally released it for certification testing in 2018. One of the main enhancements to WPA3 over WPA2 is the use of Simultaneous Authentication of Equals, which is

discussed in more depth later. For WPA3 Enterprise implementations, the minimum security protocol strength is 192-bit. However, WPA3 also enables the use of stronger security mechanisms such as 256-bit Galois/Counter Mode Protocol (GCMP-256) as well as 384-bit Hashed Message Authentication Mode (HMAC) combined with the Secure Hash Algorithm (HMAC-SHA384). It additionally supports the Elliptic Curve Digital Signature Algorithm (ECDSA), which uses 384-bit elliptic curve encryption.

Counter-mode/CBC-MAC Protocol (CCMP)

Counter-mode/CBC-MAC protocol (CCMP) is based on the Advanced Encryption Standard (AES). It provides a stronger mechanism for securing privacy and integrity over Temporal Key Integrity Protocol (TKIP), which was previously used with WPA. An advantage to CCMP is that it utilizes 128-bit keys as well as a 48-bit initialization vector. This enhancement greatly reduces the possibility of replay attacks. One drawback to using CCMP over TKIP is that it requires additional processing power. That is why you will typically see it supported on newer hardware.

Simultaneous Authentication of Equals

Simultaneous Authentication of Equals replaces the use of preshared key in the WPA3 standard. It is defined in the 802.11s standard. The use of passwords or preshared keys has always been a concern in Wi-Fi authentication. An attacker's ability to guess or brute-force the password is a viable threat. The purpose of SAE is to address this concern and enhance the authentication method used in WPA3 Personal mode. It does this by utilizing a mechanism derived from the Dragonfly Key Exchange defined in RFC 7664. For a more detailed understanding of how SAE works, see the IEEE.org resource at https://ieeexplore.ieee.org/document/4622764.

Wireless Cryptographic Protocol Summary

As you can see from the previous section, the Wi-Fi protocols and cryptographic protections continue to evolve. This is a trend that is expected in security. As computing power is enhanced, the viability to break strong encryption protocols becomes easier. Because of this, we will continue to see this evolution of cryptography and new wireless protocols. Figure 20-1 shows the latest wireless network configuration options on a macOS X system. As you can see, less secure protocols such as WEP and WPA are still available options. These options typically are made available for backward compatibility; however, it is always best to utilize the latest and most secure protocol that your device and network support.

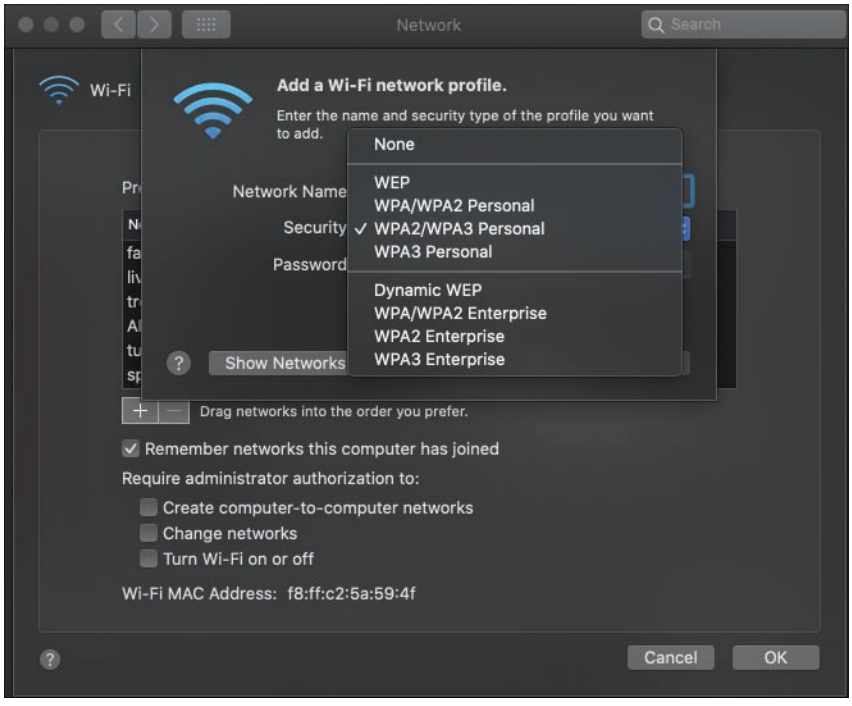


FIGURE 20-1 Wireless Client Configuration

Authentication Protocols

There are several types of technologies for authenticating a user to a local-area network. Examples that are software-based include LDAP and Kerberos, whereas an example that includes physical characteristics would be 802.1X. Keep in mind that there is a gray area between localized and remote authentication technologies. Here, each technology is placed in the category in which it is most commonly used.

In the following sections, we mention several encryption concepts that work with the various authentication technologies.

802.1X and EAP

802.1X is an IEEE standard that defines port-based network access control (PNAC). Not to be confused with 802.11x WLAN standards, ***IEEE 802.1X*** is a data link layer authentication technology used to connect hosts to a LAN or WLAN. 802.1X allows you to apply a security control that ties physical ports to end-device MAC addresses and prevents additional devices from being connected to the network. It is

a good way of implementing port security, much better than simply setting up MAC filtering.

Setting up security control starts with the central connecting device such as a switch or wireless access point. These devices must first enable 802.1X connections; they must have the 802.1X protocol (and supporting protocols) installed. Vendors that offer 802.1X-compliant devices (for example, switches and wireless access points) include Cisco, Symbol Technologies, and Intel. Next, the client computer needs to have an operating system, or additional software, that supports 802.1X. The client computer is known as the supplicant. All recent Windows versions support 802.1X. macOS offers support as well, and Linux computers can use Open1X to enable client access to networks that require 802.1X authentication.

802.1X encapsulates the ***Extensible Authentication Protocol (EAP)*** over wired or wireless connections. EAP is not an authentication mechanism in itself but instead defines message formats. 802.1X is the authentication mechanism and defines how EAP is encapsulated within messages. Figure 20-2 shows an 802.1X-enabled network adapter. In the figure, you can see that the box for enabling 802.1X has been checked and that the type of network authentication method for 802.1X is EAP—specifically, Protected Extensible Authentication Protocol (PEAP).

Key Topic

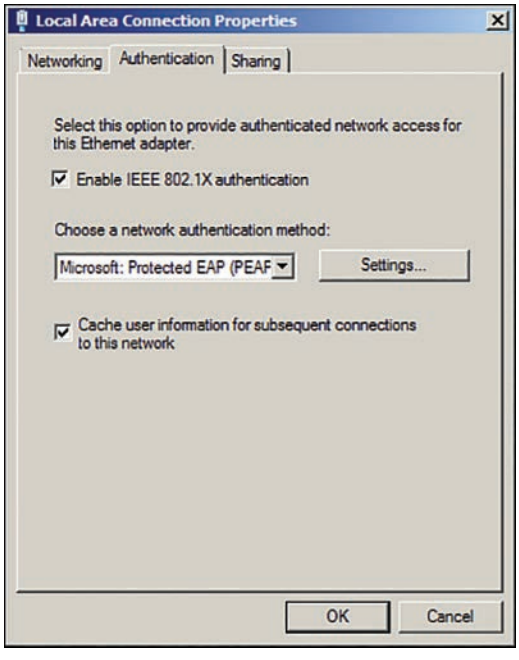


FIGURE 20-2 An 802.1X-Enabled Network Adapter in Windows

NOTE To enable 802.1X in Windows, you access the Local Area Connection Properties page.

Following are three components to an 802.1X connection:

- **Supplicant:** A software client running on a workstation. This is also known as an authentication agent.
- **Authenticator:** A wireless access point or switch.
- **Authentication server:** An authentication database, most likely a RADIUS server.

The typical 802.1X authentication procedure has four steps. The components used in these steps are illustrated in Figure 20-3.

Key Topic

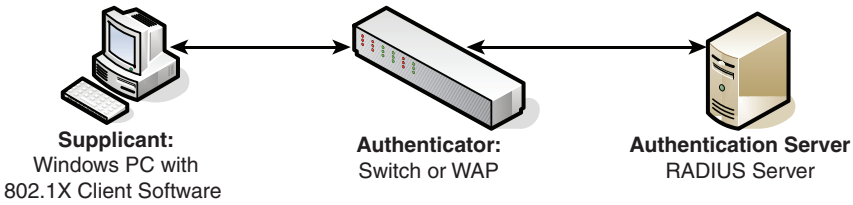


FIGURE 20-3 Components of a Typical 802.1X Authentication Procedure

Step 1. **Initialization:** If a switch or wireless access point detects a new supplicant, the port connection enables port 802.1X traffic; other types of traffic are dropped.

Step 2. **Initiation:** The authenticator (switch or wireless access point) periodically sends EAP requests to a MAC address on the network. The supplicant listens for this address and sends an EAP response that might include a user ID or other similar information. The authenticator encapsulates this response and sends it to the authentication server.

Step 3. **Negotiation:** The authentication server then sends a reply to the authenticator. The authentication server specifies which EAP method to use. (These are listed next.) Then the authenticator transmits that request to the supplicant.

Step 4. **Authentication:** If the supplicant and authentication server agree on an EAP method, the two transmit until there is either success or failure to authenticate the supplicant computer.

Following are several types of EAP authentication:

- **EAP-MD5:** This is a challenge-based authentication providing basic EAP support. It enables only one-way authentication and not mutual authentication.
- ***EAP-TLS*:** This version uses Transport Layer Security, which is a certificate-based system that does enable mutual authentication. It does not work well in enterprise scenarios because certificates must be configured or managed on the client side and server side.
- ***EAP-TTLS*:** This version is Tunneled Transport Layer Security and is basically the same as TLS except that it is done through an encrypted channel, and it requires only server-side certificates.
- ***EAP-FAST*:** This uses a protected access credential instead of a certificate to achieve mutual authentication. FAST stands for Flexible Authentication via Secure Tunneling.
- **PEAP:** The ***Protected Extensible Authentication Protocol (PEAP)*** uses MS-CHAPv2, which supports authentication via Microsoft Active Directory databases. It competes with EAP-TTLS and includes legacy password-based protocols. It creates a TLS tunnel by acquiring a public key infrastructure (PKI) certificate from a server known as a certificate authority (CA). The TLS tunnel protects user authentication much like EAP-TTLS.

IEEE 802.1x

Although 802.1X is often used for port-based network access control on the LAN, especially VLANs, it can also be used with VPNs as a way of remote authentication. Central connecting devices such as switches and wireless access points remain the same, but on the client side 802.1X would need to be configured on a VPN adapter instead of a network adapter.

Many vendors, such as Intel and Cisco, refer to 802.1X with a lowercase *x*; however, the IEEE displays this on its website with an uppercase *X*, as does the IETF. The protocol was originally defined in 2001 (802.1X-2001) and then redefined in 2004 and 2010 (802.1X-2004 and 802.1X-2010, respectively).

Remote Authentication Dial-In User Service (RADIUS) Federation

The RADIUS protocol is covered in more detail in Chapter 24, "Implementing Authentication and Authorization Solutions." This section focuses specifically on a ***Remote Authentication Dial-In User Service (RADIUS) federation***, which is used when an organization has multiple RADIUS servers—possibly on different

networks—that need to communicate with each other in a safe way. Communication is accomplished by creating trust relationships and developing a core to manage those relationships as well as the routing of authentication requests. It is often implemented in conjunction with 802.1X. This federated network authentication could also span between multiple organizations.

Methods

WEP is the weakest type of encryption. WPA is stronger, and WPA2 is the strongest of the three. However, it is better to have WEP as opposed to nothing. If this is the case, you should use encryption keys that are difficult to guess and should consider changing those keys often. Some devices can be updated to support WPA, whether through a firmware upgrade or through the use of a software add-on. Figure 20-4 shows a typical WAP with WPA2 and AES configured; AES is the cipher type.

Key Topic

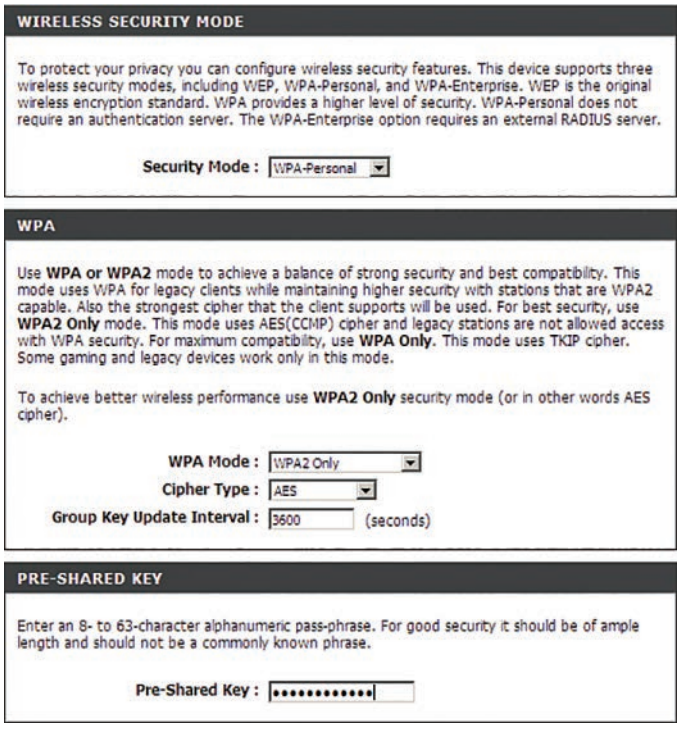


FIGURE 20-4 Wireless Security Configuration on a Typical Access Point

The ***preshared key (PSK)*** used to enable connectivity between wireless clients and the WAP is a complex passphrase. PSK is automatically used when you select WPA2-Personal in the Security Mode section. The other option is

WPA2-Enterprise, which uses a RADIUS server in this WAP. So, if you ever see the term *WPA2-PSK*, this means that the WAP is set up to use the WPA2 protocol with a preshared key, and not an external authentication method such as RADIUS.

Open Wi-Fi networks are considered insecure and require a protection mechanism be used to supplement them. VPN connections are meant to be secure sessions accomplished through an encrypted tunnel. They are best secured in a wired environment, but sometimes a wireless VPN connection is required. Some devices offer this connection but in an inherently insecure manner; this is known as VPN over *open* wireless, with "open" being the operative word, meaning insecure and unencrypted. In most business scenarios, this type of connection is unacceptable, and should be scanned for with a wireless scanning utility. Just the presence of a VPN is not enough; some kind of encryption is necessary, whether it be Point-to-Point Tunneling Protocol (PPTP), IPsec, or another secure protocol.

For example, in a standard Cisco wireless VPN configuration, the wireless client initiates a tunnel to a VPN server (a Cisco router), but it is done in a pass-through manner via a wireless LAN (WLAN) controller. (A Lightweight Access Point is often also part of the solution.) It's the router that must be set up properly, and in this scenario IPsec should be installed and configured. That enables the encryption of session data between the wireless client and WLAN controller.

In other scenarios, especially in smaller offices/home offices, a single device will act as an all-in-one solution. Though wireless VPN connections are uncommon in SOHO environments, this solution presents only a single layer of defense, and you can easily forget to initiate the proper encryption. There are many authentication mechanisms and possibilities, and several ways to encrypt the session. The key here is to remember to have clients authenticate in a secure manner and to handshake on an encryption protocol that will protect the data.

The Wireless Transport Layer Security (WTLS) protocol is part of the Wireless Application Protocol (WAP) stack used by mobile devices. It enables secure user sessions—for instance, banking transactions—using algorithms such as RSA, ECC, Triple DES, and MD5 or SHA.

Wi-Fi Protected Setup

Wi-Fi Protected Setup (WPS) is a security vulnerability. Although it was created originally to give users easy connectivity to a wireless access point, later all major manufacturers suggested that it be disabled (if possible). In a nutshell, the problem with WPS was the eight-digit code. It effectively worked as two separate smaller codes that collectively could be broken by a brute-force attack within hours.

There isn't much that can be done to prevent the problem other than disabling WPS altogether in the WAP's firmware interface or, if WPS can't be disabled,

upgrading to a newer device. In summary, WPS is a deprecated and insecure technology that should not be allowed on a wireless network.

Captive Portals

Have you ever stayed at a hotel or gone to a coffee shop that had free Wi-Fi? What happens when you use that wireless network? Chances are you are redirected to a web page that asks for authentication prior to normal Internet use. Quite often you will have to create an account with a username (usually an email address) and password, which is authenticated through email. This is an example of a captive portal. So, the captive portal method forces the HTTP client (for instance, a web browser) of the wireless device to authenticate itself via a web page. The redirection could occur as HTTP or as DNS. Quite often, it is done through basic TLS-secured HTTPS web pages. This process can often be circumvented with the use of a packet sniffer such as Wireshark. To avoid this potential hazard, an organization can opt for extended or multifactor authentication. Many free, one-time charge, and subscription-based applications are available for an organization to use with Windows and Linux-based platforms. The whole point of the technology is to be able to track users who access the free wireless network. If a user performs any suspect actions, that user can be traced by way of email address, IP address, and MAC address, in addition to other means if multifactor authentication is used.

Installation Considerations

Key Topic

Strategic ***wireless access point (WAP) placement*** is vital. That is why it is essential to perform a ***site survey*** before deploying wireless equipment. A site survey is typically performed using a variety of ***Wi-Fi Analyzer*** tools to produce a ***heat map*** of all wireless activity in the area. This map will help determine the best placement for access points. Usually, the best place for a WAP is in the center of a building. This way, equal access can be given to everyone on the perimeter of the organization's property, and there is the least chance of the signal bleeding over to other organizations. If needed, you can attempt to reduce the transmission power levels of the antenna, which can reduce the broadcast range of the WAP. Also, to avoid interference in the form of EMI or RFI, you should keep WAPs away from any electrical panels, cables, devices, motors, or other pieces of equipment that might give off an electromagnetic field. If necessary, you can shield a device creating the EM field or shield the access point itself.

To really know how to best arrange and secure your wireless connections, you need to understand the different wireless systems and antenna types available. The most common wireless system is point-to-multipoint. This system is commonly used in WLANs where a single central device (such as a SOHO wireless router) will

connect to multiple other wireless devices that could be located in any direction. Specifically, it makes use of omnidirectional antennas such as vertical omnis, ceiling domes, and so on. A typical wireless router might have two, three, four, or more vertical omnidirectional antennas. For example, with the introduction of 802.11ax, otherwise known as Wi-Fi 6, we see wireless access points and routers utilizing up to eight antennas. This is due in part to the fact that 802.11ax works on both 2.4-GHz and 5-GHz frequency bands.

With the constant innovation in wireless technology, we also see the introduction of technologies such as Multi-User Multiple Input (MU-MIMO), which utilizes a mechanism called *beamforming* for transmitting and receiving signals. This helps greatly increase the efficiency of the range. Although the technology continues to evolve, the placement of the WAPs and deployment of the antennas is still similar. Antennas can be rotated so that they are parallel to each other, or at an angle to each other; for example, 180 degrees is often a good configuration to, in essence, "sweep" the area for wireless transmissions. However, you might choose a different method. For example, you might have 100 computers on one floor and two WAPs to work with. The best method might be to position them at vertical angles from each other. One would be in the building's northeast corner and the other in the southwest corner. Then, each set of three antennas could be positioned in a 90-degree sweep as shown in Figure 20-5. As long as the building isn't larger than the range of the antennas, this setup should allow for excellent wireless coverage.

Key Topic



FIGURE 20-5 Wireless Point-to-Multipoint Layout

Addressing channel overlaps by channel selection and channel width selection can impact performance and security as well. In this case, 5-GHz frequency bands usually offer better performance than 2.4-GHz bands, but both can be monitored over the air by attackers. Some WAPs can be set to autoconfigure, seeking the least-used wireless frequencies/channels, which can be great for performance but might be a security risk. Consider using less common channel numbers and perhaps a narrower

channel—for example, 40 MHz instead of 80 MHz for 802.11ac. And remember to reduce antenna power levels as much as possible. Once again, you must test carefully, and then balance performance and security for optimal organization efficiency.

As we discussed in the previous section, 802.11ax is one of the most innovative wireless technologies to date. It is quickly becoming the most widely used for corporate and home use cases. One of the major advantages to 802.11ax is its increased efficiency and speed. However, it also has made enhancements in addressing channel congestion utilizing a new channel-sharing technique that takes advantage of a technology called orthogonal frequency-division multiple access (OFDMA). For more details about this technology, visit www.cisco.com/c/en/us/products/collateral/wireless/white-paper-c11-740788.html.

In some more simplistic point-to-point wireless systems, only two points need to be connected; these points are usually fixed in location. In this case you would use directional antennas—for example, a parabolic antenna (dish) or a Yagi antenna.

Whatever you implement, it's a good idea to perform a wireless site survey. There are three kinds of wireless surveys you might perform, each with its own purpose, and all of which are usually based on software that collects WLAN data and signal information (though hardware can also be used to measure radio frequencies):

- A passive site survey listens to WLAN traffic and measures signal strength.
- An active survey actually sends and receives data to measure data transfer rate, packet loss, and so on.
- A predictive survey is a simulated survey based on real data such as the WAP to be used, the distance between the average computer and WAP, and so on.

Surveys can be instrumental in uncovering nonaggressive interference such as neighboring radio waves and electrical equipment (RFI and EMI, mentioned earlier).

One of the tools used most commonly in site surveys is a Wi-Fi Analyzer. Wi-Fi Analyzers come in many forms. Some are hardware devices that are purpose built to analyze the wireless signals in an area. Or they can simply be a piece of software that runs on a laptop and works in conjunction with the wireless adapter either built into the laptop or a third-party adapter that connects via USB. The main purpose of the Wi-Fi Analyzer is to give you a picture of the signal and channel saturation in the area you are surveying. From a security perspective, you can also use it to create a heat map of where your Wi-Fi network signals are able to reach.

Surveys can also be used to locate aggressive *jamming* techniques often caused by wireless signal jammers. A signal jammer can easily be purchased online and can be used to initiate a denial-of-service attack to the wireless network. This attack is

made by creating random noise on the channel used by the WAP or by attempting to short out the device with powerful radio signals. The right wireless software (such as inSSIDer Paessler PRTG) can be used to locate signal jammers in or around your building so that you can remove them. Wireless software can also be used to identify potential wireless replay attacks that might exist within the network infrastructure.

Many WAPs come with a built-in firewall. If the firewall is utilized, the stateful packet inspection (SPI) option and NAT filtering should be enabled. The WAP might also have the capability to be configured for MAC filtering (a basic form of network access control), which can filter out which computers can access the wireless network. The WAP does this by consulting a list of MAC addresses that have been previously entered. Only the network adapters with those corresponding MAC addresses can connect; no one else can join the wireless network. In some cases, a device might broadcast this MAC table. If this is the case, look for an update for the firmware of the WAP and, again, attempt to fine-tune the broadcast range of the device so that it does not leak out to other organizations. Because MAC filtering and a disabled SSID can be easily circumvented using a network sniffer, it is very important to also use strong encryption and possibly consider other types of network access control (such as 802.1X) and external authentication methods (such as RADIUS).

Some WAPs also support isolation. AP isolation (also known as isolation mode) means that each client connected to the WAP will not be able to communicate with any other clients connected to the WAP. Each client can still access the Internet (or other network that the WAP is connected to), but every wireless user will be segmented from the other wireless users.

It is also possible to include the IEEE 802.1X standard for port-based network access control that can provide for strong authentication. For a WAP to incorporate this kind of technology, it must also act as a router, which adds the duty of wireless gateway to the WAP.

Another option is to consider encryption technologies on the application layer, such as TLS, SSH, or PGP; these and others can help secure data transmissions from attackers who have already gained access to the wireless network. When it comes down to it, authentication and a strong wireless protocol such as WPA3 with AES are the two security precautions that will best help protect against network attacks.

Controller and Access Point Security

We previously mentioned wireless LAN controllers. If your organization has or needs multiple WAPs, a smart method is to wire them each to a WLAN controller. This device acts as a switch for all the WAPs, thus increasing data transfer speeds between them, and more importantly, centralizes the management of security

options. Most of the WAP functionality is handed off to the controller. If this controller is physically stored in a server room or data center, the access to its functionality becomes inherently more secure (especially if a Faraday cage is implemented). In this scenario, the WAP devices are considered "thin access points" because of the reduced functionality. Contrast this to the typical "fat access point" that contains all functionality. Because the fat access point must be located close to users to be functional, it creates more potential for attackers to gain access to it. Most midsized companies and virtually all enterprise networks use WLAN controllers.

Finally, another option is to not run wireless at all. It's tough to hack into a wireless network that doesn't exist! This solution is typically seen only in very high security facilities such as government or the Department of Defense. However, if you decide to go down the antiwireless road, make sure that any devices that enable wireless access have those wireless functions disabled. This includes wireless access points, laptops, and other mobile devices that have wireless adapters, and any Bluetooth, infrared, or other wireless transmitters. The truth is, in today's world of IoT devices, it is almost impossible to go without a wireless network. Although some IoT devices do have Ethernet ports, many do not offer one at all. The best approach to integrating IoT devices into your network is to segment them from mission-critical devices. It is well known and understood that many IoT devices do not have the most robust security built in and can increase the risk of compromise on your network.

Wireless Access Point Vulnerabilities

The wireless access point is the central connecting device for wireless network adapters that might exist in PCs, laptops, handheld computers, mobile devices, and other computers. You need to secure any broadcasts that the WAP might make and verify that transmissions are encrypted with a strong encryption technique. It's also important to watch out for rogue access points and round up any nomads on your network. Many wireless controllers can identify rogue access points that are in range and enable you to label them as friendly or malicious. Obviously, a friendly access point would be one that you know of, such as a neighbor or another company in your general area. However, some access points actually may be compromised devices that are being used to deauthorize your clients and spoof your network. This type of attack would allow the rogue access point to become a man in the middle and gather credentials as well as sniff traffic on your network. By labeling the rogue unknown access points as malicious, you are able to protect your clients from connecting to them. However, it is important that you do not label a "friendly" access point as malicious. You might just upset a neighbor that way.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 20-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 20-2 Key Topics for Chapter 20

Key Topic Element	Description	Page Number
Paragraph	Cryptographic protocols	551
Figure 20-2	An 802.1X-enabled network adapter in Windows	554
Figure 20-3	Components of a typical 802.1X authentication procedure	555
Figure 20-4	Wireless security configuration on a typical access point	557
Paragraph	Installation considerations	559
Figure 20-5	Wireless point-to-multipoint layout	560

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

Wi-Fi Protected Access 2 (WPA2), Wi-Fi Protected Access 3 (WPA3), Simultaneous Authentication of Equals (SAE), IEEE 802.1X, Extensible Authentication Protocol (EAP), EAP-TLS, EAP-TTLS, EAP-FAST, Protected Extensible Authentication Protocol (PEAP), Remote Authentication Dial-In User Service (RADIUS) federation, preshared key (PSK), Wi-Fi Protected Setup (WPS), wireless access point (WAP) placement, site survey, Wi-Fi Analyzer, heat map

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is one of the key enhancements in WPA3 and a replacement for PSKl?
2. What type of EAP authentication uses the protected access credential?
3. What encryption protocol addresses some of the vulnerabilities of TKIP?
4. Which component of 802.1X is a software client?
5. Which component of 802.1X is an access point or switch?
6. Which tool can be used to get a visual picture of Wi-Fi channel saturation?
7. Which encryption protocol is used with WPA2 and WPA3?



This chapter covers the following topics related to Objective 3.5 (Given a scenario, implement secure mobile solutions) of the CompTIA Security+ SY0-601 certification exam:

- Connection methods and receivers
  - Cellular
  - WiFi
  - Bluetooth
  - NFC
  - Infrared
  - USB
  - Point-to-point
  - Point-to-multipoint
  - Global Positioning System (GPS)
  - RFID
- Mobile device management (MDM)
  - Application management
  - Content management
  - Remote wipe
  - Geofencing
  - Geolocation
  - Screen locks
  - Push notifications
  - Passwords and PINs
  - Biometrics
  - Context-aware authentication
  - Containerization
  - Storage segmentation
  - Full device encryption
- Mobile devices
  - MicroSD hardware security module (HSM)
  - MDM/Unified Endpoint Management (UEM)
  - Mobile application management (MAM)
  - SEAndroid
- Enforcement and monitoring of:
  - Third-party application stores
  - Rooting/jailbreaking
  - Sideloading
  - Custom firmware
  - Carrier unlocking
  - Firmware over-the-air (OTA) updates
  - Camera use
  - SMS/Multimedia Messaging Service (MMS)/Rich communication services (RCS)
  - External media
  - USB On-The-Go (USB OTG)
  - Recording microphone
  - GPS tagging
  - WiFi direct/ad hoc
  - Tethering
  - Hotspot
  - Payment methods
- Deployment models
  - Bring your own device (BYOD)
  - Corporate-owned personally enabled (COPE)
  - Choose your own device (CYOD)
  - Corporate-owned
  - Virtual desktop infrastructure (VDI)

Implementing Secure Mobile Solutions

This chapter starts by exploring connection methods and receivers. It then covers mobile device management (MDM) and the features and functions of this type of control. From there, the chapter moves into a discussion of security mechanisms in mobile devices such as HSM, UEM, MAM, and SEAndroid. This chapter also covers the enforcement of a number of mobile device security concerns, such as jailbreaking or rooting a device as well as sideloading and external media use, to name a few. The chapter concludes with an overview of mobile device deployment models.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 21-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 21-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Connection Methods and Receivers	1, 2, 6
Mobile Device Management	4, 5
Mobile Device Management Enforcement and Monitoring	3
Mobile Devices	7
Deployment Models	8

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following has benefits for contactless payments?
   a. NFC
   b. USB
   c. LTE
   d. None of these answers are correct.

2. Which of the following is utilized to alert when users enter or exit an organization's physical borders?
   a. Point-to-point deployment
   b. Geofencing
   c. Infrared
   d. All of these answers are correct.

3. Which is a way to reduce the impact of mobile device theft?
   a. Full device encryption
   b. Application management
   c. Context-aware authentication
   d. None of these answers are correct.

4. Which MDM feature allows software updates to be pushed to devices remotely?
   a. GPS
   b. OTG
   c. WAP
   d. OTA

5. Which of the following terms means loading third-party apps from a location outside the official application store for that device?
   a. Carrier unlocking
   b. SIM card cloning
   c. Jailbreaking
   d. Sideloading

6. Which of the following terms means sending unsolicited messages to Bluetooth-enabled devices?
   a. Bluesnarfing
   b. Bluejacking
   c. Bluehacking
   d. None of these answers are correct.

7. Which of the following is one of the best ways to ensure that data is secured and that applications work properly without interference from potential attackers?
   a. Application whitelisting
   b. Encryption
   c. Application management
   d. None of these answers are correct.

8. Which of the following describes a policy in which the company supplies employees with phones?
   a. BYOD
   b. COPE
   c. VDI
   d. None of these answers are correct.

Foundation Topics

Key Topic

Connection Methods and Receivers

Here we cover the various connection methods and receivers that are included with most modern mobile devices today. We also cover some of the security concerns with these technologies and finish with an overview of secure implementation best practices.

Bluetooth, radio frequency identification (RFID), and near-field communication (NFC) are not wireless networking technologies in the general sense the way Wi-Fi is. But anything that has two or more wireless devices that communicate with each other could technically be considered a wireless network.

Like any wireless technology, Bluetooth is vulnerable to attack. Bluejacking and bluesnarfing are two types of vulnerabilities to Bluetooth-enabled devices. Bluetooth is also vulnerable to conflicts with other wireless technologies. For example, some WLAN (or Wi-Fi) standards use the 2.4-GHz frequency range, as does Bluetooth, and even though Bluetooth uses frequency hopping, conflicts can occur between 802.11g or 802.11b networks and Bluetooth personal area networks (PANs). To avoid this, you should use Bluetooth version 1.2 devices or greater, which employ adaptive frequency hopping, improving resistance to radio interference. Also, you should consider placing Bluetooth access points (if they are used) and WLAN access points in different areas of the building. Some companies have policies governing Bluetooth usage; in some cases, it is not allowed if 802.11 standards are in place, and in some cases, a company may enforce rules that say Bluetooth can be used only outside the building. In other cases, a company may put its 802.11 devices on specific channels or use WLAN standards that use the 5-GHz range.

Bluetooth-equipped devices can use ***near-field communication (NFC)***, which allows two mobile devices (or a mobile device and a stationary computer) to be automatically paired and transmit data. NFC is not limited to Bluetooth, but Bluetooth is probably the most common technology used to transmit data wirelessly over short distances. Of course, even though the distance is short, attackers can still eavesdrop on it. In addition, NFC is a data transmission protocol, but not necessarily secure. Data can be destroyed by use of a jammer, and users are also at risk of replay attacks. As of this writing, NFC does not offer preventive security in this respect, but users can prevent these attacks by using only applications that offer SSL/TLS or other secure channels during an NFC session.

Anytime a cell phone or smartphone connects, it uses some type of wireless service. Whether it's 5G, 4G, 3G, GSM, Wi-Fi, infrared, RFID, or Bluetooth, security implications exist. To minimize risks, the best solution is to turn off the particular

service when not in use, use airplane mode, or simply turn off the mobile device altogether if it is not being used.

Bluetooth is especially vulnerable to virus attacks, as well as bluejacking and bluesnarfing. *Bluejacking* is the sending of unsolicited messages to Bluetooth-enabled devices such as mobile phones. You can stop bluejacking by setting the affected Bluetooth device to undiscoverable or by turning off Bluetooth altogether.

Bluesnarfing is the unauthorized access of information from a wireless device through a Bluetooth connection. Generally, bluesnarfing is the theft of data (calendar information, phonebook contacts, and so on). Ways of discouraging bluesnarfing include using a pairing key that is not easy to guess; for example, you should stay away from 0000 or similar default Bluetooth pairing keys! Otherwise, you should set Bluetooth devices to undiscoverable (only after setting up legitimate Bluetooth devices, of course) or turn off Bluetooth altogether.

Wi-Fi has many vulnerabilities as well. Not only should mobile devices connect in a secure, encrypted fashion, but also you, as security administrator, need to keep a sharp eye on the current Common Vulnerabilities and Exposures (CVEs) and the available updates and patches for those vulnerabilities. For example, there was a flaw in the programming of a well-known Wi-Fi System on Chip (SoC). The firmware had a vulnerability that could result in buffer overflows, which could then be exploited by attackers—connecting remotely via Wi-Fi—ultimately enabling the execution of their own code. Sometimes SoCs are not properly vetted for vulnerabilities, so you must be ready to patch at a moment's notice. This approach applies not only to smartphones and other typical mobile devices but also to just about all devices in the Internet of Things (IoT) that have built-in Wi-Fi connections.

RFID and NFC

Radio frequency identification (RFID) has many uses, but it all boils down to identifying and tracking tags that are attached to objects. In the security world, that generally means authentication.

As with any wireless technology, RFID is susceptible to attack. For example, some RFID tags can be affected by skimming, on-path attacks (formerly known as man-in-the-middle, or MITM, attacks), sniffing, eavesdropping/replaying, spoofing, and jamming (DoS). From an authentication standpoint, attackers are using these attacks to try to find out the passcode. An RFID tag can also be reverse engineered if attackers get possession of it. Finally, power levels can be analyzed to find out passwords. On some RFID tags, correct passcodes emit a different level of power than incorrect passcodes. To prevent these attacks, you (or your security admin team) should consider newer generation RFID devices, encryption, chip coatings, filtering of data, and multifactor authentication methods. Encryption is one of the best methods.

Included in this prevention method are rolling codes, which are generated with a pseudorandom number generator (PRNG), and challenge-response authentication (CRA), where the user (or user's device) must present the valid response to the challenge.

RFID ranges vary depending on the EM band used—from 10 cm up to 200 meters. Many authentication readers can be as much as 1 meter away, which is enough to facilitate skimming of information by attackers. One way to avoid skimming is to use newer RFID proximity readers and keys—ones that use lower frequencies—from respectable companies. Another way is to utilize the set of protocols called near-field communication. NFC generally requires that communicating devices be within 4 cm of each other, which makes skimming of information difficult. If an employee uses a smartphone to enter a building instead of an RFID device, NFC should be implemented. NFC has obvious benefits for contactless payment systems or any other non-contact-oriented communications between devices. However, for optimal security, you should use contact-oriented readers and cards.

More Wireless Connection Methods and Receivers

Let's not forget about cellular connections. Many companies don't allow cellular access, meaning 2G, 3G, 4G, 5G LTE, and so on. These connections are often denied within company premises, and instead the company relies on Wi-Fi methods—for example, Wi-Fi calling. This is common in choose-your-own-device (CYOD) and corporate-owned, personally enabled (COPE) environments. But if cellular is necessary in a bring-your-own-device (BYOD) environment, security can be increased by using newer devices, updating the devices' operating systems, updating the preferred roaming list (PRL), updating identification technologies (such as IMEI and IMSI), and using a VPN for data connections over cellular. You should use these methods for employees who must be on the road as well. And if there are foreseeable instances where cellular is not available for mobile employees, you should be sure that they understand the risks of open Wi-Fi networks and that they should avoid them as much as possible.

When it comes down to it, the use of a ***Global Positioning System (GPS)*** in general should be examined carefully, weighing the benefits against the possible vulnerabilities. This includes GPS derivatives such as GPS tagging, geofencing, and geolocation. Many executives and other employees use their mobile devices at work, which brings up many security concerns besides GPS. Collectively, they are known as BYOD concerns and are described in the following sections.

For the purposes of privacy, it is best to disable GPS whenever possible. Also, you should consider disabling other GPS and geolocation-related technologies. For example, in *geotagging*, geographical identification information, such as latitude

and longitude coordinates, is added to photographs, websites, SMS messages, and more. It is common in social media and can be a great tool, but it can also be an easy way for attackers to zero in on high-profile executives and other employees of an organization. In these cases, you should consider a *geofence*—a virtual fence defining the boundaries of an actual geographical area. Geofencing is an excellent way to be alerted to users entering and exiting an organization's physical premises and can provide security for wireless networks by defining the physical borders and allowing or disallowing access based on the physical location of users, or more accurately, the users' computers or mobile devices!

Some organizations rely on satellite communications (SATCOM): sometimes for long-distance communications and sometimes for communicating between buildings in a campus in a point-to-point or point-to-multipoint deployment. Either way, it is important to understand that SATCOM devices can be at risk if their firmware is not updated. Exploits could include the installation of malicious firmware and the execution of arbitrary code. Updating may require physical access to the parabolic antenna using a laptop or terminal device. But remember, the easier it is to access these antennas, the more likely they can be hacked by malicious individuals. Of course, in some cases these antennas can be hacked remotely as well. Secure planning is necessary when it comes to physical access, firewalling (if possible), and especially updating.

A wireless connection method that is often overlooked when it comes to security is infrared. It is employed in many technologies that we use every day, including mobile phones. Consequently, it is often not secured properly. Of course, many wireless receivers can be added to a device by simply plugging in a USB adapter. This capability is also often overlooked.

Remember this: wireless technologies are always evolving. But anything is hackable, given time and ingenuity. When it comes to over-the-air (OTA) technology, you should be ever vigilant: know the latest exploits and prepare a worthy defense.

Secure Implementation Best Practices

When it comes to the best practices of secure implementation of communication methods and receivers, they are all similar. Therefore, the same controls can apply in most situations. Some of the biggest security issues with mobile device connection methods and receivers are vulnerabilities in the actual firmware and/or the software that runs on the device to interface with these technologies. These vulnerabilities can be used to interrupt or capture the communications between devices. The way to address this type of concern is to always verify that you are using the latest software version available for the device or radio. When it comes to mobile devices, typically these types of software updates are packaged into a larger update for the

device that includes other fixes and features. However, sometimes a vulnerability is so severe that it requires a hotfix or patch to quickly address the vulnerability. Either way, keeping your device up to date is a solid approach to this concern.

There is software on your device that interacts with the hardware radio. That access can be controlled through the operating system permissions. By limiting access to the use of these communication mechanisms, you are essentially reducing the attack surface. Doing so helps mitigate the impact of software vulnerabilities as they come up. Lastly, if you are not using a specific connection method, you can simply turn it off. This again reduces your attack surface. Of course, turning off a connection is not always possible, so it is always a decision you need to make based on the environment you are in. For instance, if you're at a large security conference such as DEFCON, turning off all unnecessary connection methods is a very good idea. If you are just sitting around your house, however, there is obviously less of a concern that someone might be acting in a malicious way to attack your device. We hope these best practices give you an idea of how to reduce the risk when implementing mobile solutions.

Mobile Device Management



The key to having a successful BYOD implementation is to implement storage segmentation—a clear separation of organizational and personal information, applications, and other content. It must be unmistakable where the data ownership line occurs. For networks with a lot of users, you should consider third-party offerings from companies that make use of ***mobile device management (MDM)*** platforms. These centralized software solutions can control, configure, update, and secure remote mobile devices such as Android, iOS, BlackBerry, and so on, all from one administrative console. The MDM software can be run from a server within the organization or administered within the cloud. It makes your job as a mobile IT security administrator at least manageable. From a central location, you can carry out the tasks of application management, content management, and patch management. You can also set up more secure levels of mobile device access control.

Access control is the methodology used to allow access to computer systems. For larger organizations, MDM software makes it easy for you to view inventory control, such as how many devices are active for each of the mobile operating systems used. It also makes it simpler to track assets, such as the devices themselves, and the types of data each contains.

In addition, MDM software makes it less complicated to disable unused features on multiple devices at once, thereby increasing the efficiency of the devices, reducing their footprint, and ultimately making them more secure. For instance, an employee who happens to have both a smartphone and a tablet capable of making cellular calls

doesn't necessarily need the latter. As administrator, you could disable the tablet's cellular capability, which would increase battery efficiency as well as security for that device.

User acceptance of BYOD is mixed in its reactions. Some employees like the idea of using their own device (which they might not have been allowed to use at work previously) and not having to train on a separate work computer. However, some employees believe that BYOD is just a way to move computing costs from the company to the user, and the level of trust is low. This brings up a variety of legal concerns, such as the right to privacy. Companies that offer BYOD MDM solutions counter this perception by drawing a clear line in the sand, defining exactly what employers can see (for example, corporate email) and what they can't see (such as personal texts). In general, these companies try to protect the privacy of individuals. Many organizations write clear privacy policies that define, if necessary, selective wipes of secure corporate data while protecting personal data.

Part of the debate over BYOD includes some additional concerns; for example, additional legal concerns exist about employee misconduct and fraud. Anything found that could possibly implicate an employee of wrongdoing would have to be found in the organizational portion of the data. From a forensics point of view, however, and because the device can't be split in two, if any potential wrongdoing is investigated, the device would need to be confiscated for analysis.

Most employees (of all age groups) are also concerned with how on-board devices (such as on-board cameras) can be used against them with or without their knowledge. Companies that offer BYOD solutions tend to refer to the camera (and photos/video taken) as part of the personal area of the device. However, those same companies include GPS location as something they can see, but this can be linked to a corporate login, with GPS tracking users only when they are logged in. Onboarding and offboarding in general are other concerns. Essentially, *onboarding* means that you, as security administrator, take control of the device temporarily to configure it, update it, and perhaps monitor it, and *offboarding* means that you relinquish control of the device when finished with it. This control brings up some questions for employees: When does it happen? How long does it last? How will my device be affected? Are there any architectural/infrastructural concerns? For example, will the BYOD solution change the core files of my device? Will an update done by a person when at home render the device inactive the next day at work? That's just the tip of the iceberg when it comes to questions and concerns about BYOD. The best course of action is for an organization to set firm policies about all of these topics.

Policies that need to be instituted include an acceptable use policy, data ownership policy, and support ownership policy. In essence, these policies define what users are allowed to do with the device (during work hours), who owns what data and

how that data is separated, and under what scenarios the organization takes care of technical support for the device as opposed to the users.

To help secure the mobile devices in a BYOD enterprise environment, some third-party providers offer an embedded certificate authority for managing devices and user identity, sophisticated posture monitoring and automated policy workflow so that noncompliant devices do not get enterprise access, and certificate-based security to secure email and reduce the chance of data loss.

Unfortunately, smartphones and tablets (and other mobile devices) can be the victims of attack as well. Attackers might choose to abuse your service or use devices as part of a larger-scale attack and possibly to gain access to account information. Though mobile devices can be considered computers, there are some other factors to consider specifically for mobile devices.

Users of mobile devices should be careful about giving their phone number to others and should avoid listing their phone number on any websites, especially when purchasing products. You should train your users not to follow any links sent by email or by text messages if they are unsolicited. (If there is any doubt in a user's mind, then it is best to ignore the communication.) Explain the issues with much of the downloadable software, such as games and ringtones, to your users. Also, use a locking code/password/gesture that's hard to guess; this locks the mobile device after a specific amount of time has elapsed. In addition, use complex passwords when necessary—for example, if required by company policy.

In general, mobile operating system software must be updated just like desktop computer software. You should keep these devices up to date, and there will be less chance that they will be affected by viruses and other malware. You can encrypt data in several ways; some organizations even have policies that specify how data will be encrypted. More good general tips are available at the following National Cyber Awareness System (NCAS) and U.S. Computer Emergency Readiness Team (US-CERT) website links:

www.us-cert.gov/ncas

https://us-cert.cisa.gov/ncas/tips (Go to the Mobile Devices section.)

MDM Security Feature Concerns: Application and Content Management

Let's speak more about the applications' security on mobile devices. We've already mentioned that applications should (usually) be updated to the latest version and discussed the importance of proper user interaction, but let's delve a bit deeper and talk about ways to encrypt data that is transferred through applications.

Encryption is one of the best ways to ensure that data is secured and that applications work properly without interference from potential attackers. However, you should consider whole device encryption, which encrypts the internal memory and any removable (SD) cards. Sometimes you might forget about one or the other. Then there is data in transit—data that is on the move between a client and a server. Most applications for mobile devices communicate with a server of some sort; for example, when a person uses a web browser, an email client, a contacts database, or actual "apps" that work independently of a browser but operate in a similar manner, meaning that they ultimately connect to a server. Weather apps, games, social media apps, and so on all fall into this category.

Let's consider the web browser, for instance. A mobile device connects to websites in a similar manner to a desktop computer. Basic websites use a Hypertext Transfer Protocol (HTTP) connection. But websites that require any type of personally identifiable information (PII) use HTTP Secure (HTTPS). This can then utilize one of several types of encryption, such as Transport Layer Security (TLS).

Whatever the security protocol, the important point here is that the server you are connected to makes use of a database that stores encryption keys. The key (or a portion thereof) is sent to the client device and is agreed upon (handshaking occurs) so that the transfer of data, especially private information, is encrypted and protected. Often, HTTPS pages are used to aid in the process of authentication—the confirmation of a person's (or computer's) identity, typically with the help of a username/password combination. Examples include when you log in to your account with a bank or with a shopping portal.

One of the important roles for the server is key management—the creation, storage, usage, and retirement of encryption keys. Proper key management (and the regular updating of keys) is your primary concern as security administrator. Generally, an organization purchases a master key algorithm from a third-party company such as VeriSign. That company informs the organization if a key has become compromised and needs to be revoked. These third parties might also take part in credential management (the managing of usernames, passwords, PINs, and other passcodes, usually stored within a secure database) to make things a bit easier for you. Whether this is the case depends on the size of the organization and its budget. Key management gets quite in depth, as you can imagine. For now, realize that a mobile device is an easy target. Therefore, applications (especially third-party apps) should be scrutinized to make sure they are using a solid encryption plan when personal information is transferred back and forth.

Authentication to servers and other networks (and all their applications) can get even more complicated when the concept of transitive trust is implemented. Effectively, a transitive trust occurs when two networks (or more) have a relationship such that users logging in to one network get access to data on the other. In days gone by,

these types of trusts were created automatically between different sections of networks; however, it was quickly realized that this type of transitivity was insecure, allowing users (and potential attackers) access to other networks that they shouldn't have had access to in the first place. There's a larger looming threat here as well. The *transitive trust* is based on the transitive property in mathematics, which states that if A is equal to B, and B is equal to C, then A is automatically equal to C. Put into computer terms: If the New York network trusts the California network, and the California network trusts the Hong Kong network, then the New York network automatically trusts the Hong Kong network. You can imagine the security concerns here, as well as the domino effect that could occur. So, organizations usually prefer the nontransitive trust, where users need to be authenticated to each network separately, and therefore are limited to the applications (and data) they have access to on a per-network basis.

To further restrict users and increase application security, ***allow lists*** (*application whitelisting)* are often used. This means that you, as administrator, create a list of approved applications and that users can work with only those applications, and no others. This is often done within a computer policy and can be made more manageable by utilizing a mobile device management system (which we describe elsewhere in this chapter). Users often need access to several apps: phone, email, contacts, and web browser. These applications would make up the allow list, and if a user tried to use other apps, that user would be denied, or at the very least, would be prompted for additional user credentials. If a user needed access to another app, such as the camera, you would weigh the security concerns (GPS, links to social media, and so on) and decide whether to add the app to the whitelist. Whitelisting can also be helpful when dealing with apps that utilize OAuth—a common mechanism used by social media companies to permit users to share account information with third-party applications and websites. Contrast the concept of using allow lists with ***block lists/deny lists***—the denial of individual applications—a common method used when working with email, and by antivirus and hardware-based intrusion detection systems (HIDS) programs.

MDM Security Feature Concerns: Remote Wipe, Geofencing, Geolocation, Screen Locks, Passwords and PINs, Full Device Encryption

More than 100 mobile devices end up missing (often stolen) every minute. Yes—every minute! You can imagine the variety of reasons why these thefts occur. The worst attack that can be perpetuated on a smartphone or tablet is theft. The theft of a mobile device means the possible loss of important data and personal information. There are a few ways to protect against this loss of data and to recover from the theft of a mobile device if it does happen.

First, mobile devices in an organization should utilize data encryption. The stronger the encryption, the more difficult it is for a thief to decode and use the data on the device. If at all possible, you should use full device encryption. Most modern mobile device operating systems such as Apple iOS and Android have this capability built in. So there really is no reason not to have it enabled on devices. The actual conversations on phones can also be encrypted. Voice encryption can protect the confidentiality of spoken conversations and can be implemented with a special microSD chip (preferably) or with software.

Mobile devices should also be set up for GPS tracking so that they can be tracked if they are lost or stolen. The quicker a device can be located, the less risk of data loss, especially if it is encrypted. However, GPS tracking can also be a security vulnerability for the device and possibly the user if an attacker knows how to track the phone.

The beauty of mobile devices is in their inherent portability—that and the ability to track SIM cards. If you are an administrator for mobile devices, you should consider remote *lockout* programs. If a device is lost or stolen, you can lock the device, disallowing a would-be attacker access. In addition, the device can be configured to use the "three strikes and you're out" rule, meaning that if a user tries to be authenticated to the device and is unsuccessful after three attempts, that user is locked out.

Of course, we all know that password authentication is not the strongest authentication method. Utilizing multifactor authentication is a way to enhance the strength of your authentication security—for instance, utilizing biometrics as one of the authentication factors alongside password/PIN authentication. Most modern mobile devices and even tablets have some sort of biometrics capability, such as fingerprint or face recognition. Taking it to the next level, if the data is extremely sensitive, you might want to consider a remote wipe program. If the mobile device is reported as lost or stolen, these programs can remove all data from the phone in a bit-by-bit process, making it difficult (if not impossible) to recover. This process is known as *sanitizing* the phone remotely. Of course, a solid backup strategy should be in place before a data sanitizing or remote wipe solution is implemented.

Screen locks, complex passwords, and taking care when connecting to wireless networks are also important. Though a screen lock won't deter knowledgeable attackers, it will usually deter the average person who, for example, finds a stray phone sitting in a coffee shop, mall, or other public location. User training should be implemented when users first receive their devices. Though many organizations don't take the time for training, it is a great way to show users how to secure their device, while checking whether their encryption, GPS tracking, and other features are working properly. They can also be trained on how to inform your organization and local law enforcement in the case that a device is lost or stolen, effectively reducing the risk of data loss by allowing you to find the device faster or mitigate the problem in other ways.

Another important function of an MDM solution is providing context-aware authentication by limiting or preventing access to organization resources based on the device profile and security posture (for example, a device that is rooted should not be able to access certain resources).

NOTE In the case of theft, the two best ways to protect against the loss of confidential or sensitive information are encryption and a remote wipe program.

Table 21-2 summarizes some of the mobile security concerns that are addressed with implementing mobile device management. With a mixture of user adherence to corporate policies, the workplace respecting the user's right to privacy, and a strong security plan, BYOD can be a success.

Key Topic

Table 21-2 Mobile Device Security Concerns and Countermeasures

Mobile Device Security Topic	Countermeasure
Malware	Update device to latest version (or point release for the current version). Use security suites and AV software. Enable them if preloaded on the device and update regularly. Train users to carefully screen email and selectively access websites. Be careful of social networks and third-party apps.
Botnets and DDoS	Download apps from a legitimate source. If BYOD is in place, use company-approved apps. Refrain from rooting or jailbreaking the device. Have data backed up in case the device becomes part of a botnet and has to be wiped.
SIM cloning	Use V2 and newer cards with strong encryption algorithms.
Wireless attacks	Use a strong password for the wireless network. Turn off unnecessary wireless features such as mobile hotspots and tethering. Disable Bluetooth if not in use for long periods of time (also to conserve the battery). Set the device to undiscoverable.
Theft	Utilize data and voice encryption (especially in BYOD implementations). Implement lockout, remote locator, and remote wipe programs. Limit the amount of confidential information stored on the device. Use screen locks and complex passwords.

Mobile Device Security Topic	Countermeasure
Application security	Use encryption from reputable providers. Use antimalware endpoint protection platforms. Utilize nontransitive trusts between networks and apps. Allow list (whitelist) applications. Disable geotagging.
BYOD concerns	Implement storage segmentation. Utilize an MDM solution. Create and implement clear policies that the organization and users must adhere to. Consider CYOD or COPE as opposed to the traditional BYOD method.

Mobile Device Management Enforcement and Monitoring

Mobile device management controls the deployment, operations, and monitoring of mobile devices used to access organization resources. MDM is used to enforce an organization's security policy on mobile devices. Let's look at some of the features and controls that are part of a typical MDM solution and why they are implemented:

Key Topic

- Disabling user access to third-party application stores. It is very important to control the access to application stores that are not approved by your company policy. Different application stores perform different levels of application security validation before they allow the application to be available on their store. Some of them do not perform any security validation. Therefore, the applications that are posted on these stores have a higher risk of containing malicious software.
- Restricting user or application access to mobile device hardware, such as digital cameras, network interfaces, GPS, and services or native applications such as the built-in web browser or email client.
- Rooting/jailbreaking/sideloading. Insecure user configurations such as rooting and jailbreaking can be blocked from MDM, as can ***sideloading***—the art of loading third-party apps from a location outside the official application store for that device. Note that sideloading can occur in several ways:
  - By direct Internet connection (usually disabled by default)
  - By connecting to a second mobile device via USB OTG (USB On-The-Go) or Bluetooth
  - By copying apps directly from a microSD card or via tethering to a PC or Mac

Application control becomes easier as well. Applications can be installed, uninstalled, updated, and secured from that central location. Even devices' removable storage (often USB-based) can be manipulated—as long as the removable storage is currently connected to these devices.

- Monitoring, alerting, and reporting on policy violation (for example, if a user is trying to root the mobile device). It is important to keep an eye on policy violations that might occur on a mobile device because they could lead to a compromise of the device and/or sensitive data.
- Encrypting data communication between the device and the organization as well as full device encryption of data stored on the device or in removable storage. Mobile devices often connect to nontrusted networks. This means that anyone on that network can view the data traversing the network. That is why encryption of the data communication from the device to the organization is essential. Encryption is typically set using a VPN. There are also new hybrid-like approaches such as Secure Access Service Edge (SASE) solutions. Also, because of the "mobile" aspect of these devices, there is more of a concern that such devices might be lost or stolen. If they are lost or stolen, the risk of data leakage is greatly reduced if these devices have full encryption implemented. Gleaning data from a device with full encryption and strong authentication will be very difficult.
- Providing the ability to remotely wipe the device in case the device is lost or stolen and in case of device reuse. Devices are lost or stolen every day. Most of us have had this happen to us at some point. If a device containing corporate data is lost, the organization must have the capability to remotely wipe the device to prevent access to that data.
- Enforcing strong password and PIN code authentication for accessing the device and/or organization resources. This includes password strength policies, clipping level, and so on. Enforcing authentication is usually one of the first things that is required on a device when it first connects to a corporate network and registers with the MDM. Without it, a lost device would be open to anyone who acquired it and so would the data stored on it.
- Remotely locking the device and screen and remotely resetting the password. Devices are lost or stolen every day. Most of us have had this happen to us at some point. If a device containing sensitive data is lost, the organization must have the capability to remotely lock the device to prevent access to that data.
- Enabling the enforcement of data loss prevention (DLP) on mobile devices. Mobile devices on a corporate network have the same capabilities as laptops. This means that corporate data can be easily moved to and from mobile devices. This is why it is critical to implement a policy that will enforce controls for data loss prevention.

- Performing application management by restricting the types of applications that can be installed; for example, via allow listing (whitelisting) or block list/deny listing (blacklisting) and which resources the applications can use. Due to the large threat that untrusted applications could pose to the organization, application management is usually handled within a mobile application management (MAM) framework.
- Disabling the ability for users to install custom firmware or to detect custom firmware. A jailbreak and/or root access is normally necessary to install custom firmware on a device. For this reason, it is important to validate that the device has not been tampered with in this way before allowing it on the network. With Android devices, you can also disable ADB access, for example.
- Disabling mobile device camera use. This is necessary in some highly secure environments such as government buildings. For this reason, it is important for an MDM solution to have the capability to disable camera use. There are, of course, other reasons to disable the use of a camera. This is only one sample use case.
- Pushing ***firmware over-the-air (OTA) updates*** to managed devices. Keeping a device updated is critical in maintaining the secure posture of any device. Mobile devices typically require additional control over which updates are installed and when they are installed. For instance, if a critical vulnerability that affects a specific mobile device is released, the rollout of the fixed software needs to happen very quickly. Remember, these devices are typically connected directly to the Internet all the time, so they are at a higher risk for exploitation.
- Enabling/disabling Short Message Service (SMS)/Multimedia Messaging Service (MMS)/Rich Communication Services (RCS). Of course, SMS and MMS are typically used in text messaging applications. Sometimes organizations need to disable the use of these applications permanently or temporarily to reduce the risk of data leakage as well as threats from SMS-type phishing attacks, which are sometimes called smishing attacks.
- Disabling the use of external media. The capability to disable use of external media can be useful for protecting users against possible malware that is spread via external media. It can also be used as a mitigation for data loss.
- Using USB On-The-Go (USB OTG) for attaching external devices to your mobile device via the USB connection. They can be USB flash drives, adapters, the keyboard, the mouse, and so on. Anything you physically connect to the USB port on a mobile device can possibly be malicious. For this reason, it is recommended that you disable this capability using an MDM solution.
- Deactivating the recording microphone. This feature of MDM is primarily useful when in highly secure environments.

- Disabling carrier locking or SIM cloning (also known as phone cloning), which allows two phones to utilize the same service and allows attackers to gain access to all phone data. V1 SIM cards had a weak algorithm that made SIM cloning possible (with some expertise). However, V2 cards and higher are much more difficult (if not impossible) to clone due to a stronger algorithm on the chip. Users and administrators should be aware of the version of SIM card being used and update it (or the entire smartphone) if necessary. There are techniques available to unlock a smartphone from its carrier. Users should be advised against this, and you should create and implement policies that make unlocking the SIM card difficult, if not impossible. Unlocking the phone—making it SIM-free—effectively takes it off the grid and makes it difficult to track and manage. When the SIM is wiped, the international mobile subscriber identity (IMSI) is lost, and afterward the user cannot be recognized. However, you can attempt to block list/deny list (blacklist) the smartphone through its provider using the international mobile equipment identity (IMEI), electronic serial number (ESN), or mobile equipment identifier (MEID). The ID used varies depending on the type and age of smartphone. Regardless, as a security administrator, you should avoid that tactic altogether because the damage has already been done; so, protection of the SIM becomes vital.
- Disabling GPS tagging. Most mobile devices contain a GPS. This feature is obviously useful for many reasons, such as use of maps and location services. However, it can also create a risk that users are leaking their location in files such as pictures taken on the device. When these pictures are posted on the Internet, the metadata would still have this GPS location information. For this reason, it is necessary for MDM to have the capability to disable the GPS hardware on the device.
- Enabling/disabling Wi-Fi direct/ad hoc network access. Wi-Fi direct is a method for devices to communicate directly with each other without the use of an access point. Wi-Fi ad hoc is a network of mobile devices that communicate directly with each other without the use of a network. Either can be a bad thing if implemented incorrectly. For instance, if a mobile device is configured to automatically join ad hoc networks, it might just be allowing communication from an untrusted device. These types of communication should be disabled on a corporate managed device to reduce the risk of compromise from untrusted devices.
- Enabling/disabling payment methods such as Apple Pay or Google Pay can be a requirement for some policies. If a device is corporate owned and used only for work purposes, you would want to have a way to control these methods.
- Enabling/disabling using tethering as a method for allowing another device, such as a laptop, to connect to the Internet through your phone. In general,

tethering is accomplished via USB cable. It can be useful for people who are on the road and need to access the Internet from another device. However, it can also allow data transfer between a connected device and the mobile device. This capability opens up the device to yet another attack surface. Consequently, it is important that an MDM solution have the capacity to enable or disable the tethering capability of the mobile device.

- Enabling/disabling using a hotspot to allow access to the Internet through the mobile device. This feature is similar to tethering. However, with a hotspot, instead of using a physical connection like USB to connect to the device, it utilizes wireless. Typically, this is done via Wi-Fi but can also be done using Bluetooth connectivity. When you enable this feature, you are essentially turning your mobile device into a Wi-Fi access point. If not properly secured, this hotspot can be used by anyone within range. This capability, of course, opens up an attack surface, but it also runs the risk of high data usage. For this reason, many organizations disable the hotspot feature using an MDM solution.

Mobile device management capabilities could be offered by the mobile vendor or provided by a third-party management tool that offers multivendor support. The second option is currently the most used due to the increased adoption of BYOD and heterogeneous types of devices used within an organization.

One of the characteristics of an MDM solution is the use of over-the-air device management. OTA historically refers to the deployment and configuration performed via a messaging service, such as Short Message Service (SMS), Multimedia Messaging Service (MMS), Rich Communication Services (RCS), or Wireless Application Protocol (WAP). Nowadays it's used to indicate remote configuration and deployment of mobile devices, including the following benefits/needs:

- Higher level of control
- Intellectual property retention
- Regulatory compliance (for example, if it is not possible to store data in the cloud)



Mobile Devices

A ***mobile application management (MAM)*** framework can be used to address the types of applications that can be installed (for example, by using allow lists or block/deny lists and which resources the applications can use). Due to the large threat that untrusted applications may pose to an organization, application management is usually handled within the MAM framework.

The messaging app is a particularly devious gateway for attackers. SMS, MMS, and RCS are vulnerable to malware, and unwary users of mobile devices are especially susceptible to Trojans and phishing via SMS texts. One way to prevent these types of attacks is to install mobile antimalware in the form of one mobile security suite or another. This endpoint protection platform needs to be updated and is best controlled from an MDM solution. Another way is to block messaging apps altogether or to use company-approved messaging apps. This option depends on what type of mobile environment you allow. It works for some mobile environments where the IT department has more control over devices, but probably not for BYOD.

If your organization uses a mobile payment method, it is important to understand that the applications that control these payment methods and the devices they run on can have several vulnerabilities. They include weak passwords (for the mobile device and for the payment app), user error, and phishing. Not only that, the technology itself is inherently insecure given the mobility of the devices. Users should be educated about not using their mobile devices for payment while making a public Wi-Fi connection. They should also be taught how to properly and discreetly use payment apps. And, of course, they should be instructed on how to avoid loss or theft of mobile devices and what to do if loss or theft occurs. As the security administrator, you should consider using an external reader for payment transactions on mobile devices and teaching users to keep the reader separate from mobile devices when not in use.

Geotagging (also written as *geo-tagging*) is another application concern. Photos, videos, websites, messages, and much more can be geotagged. Geotagging is the adding of data to the content in question, helping users to gather location-specific information. For example, if a user wanted to take a picture of a favorite store at the mall and help friends find it, that user could geotag the picture. However, doing so requires that the smartphone (or other mobile device) have GPS installed and running. This then means that the user's smartphone can be physically located and tracked. Depending on the applications running, being tracked could pose a security threat. In a corporate environment, the security administrator often chooses to disable geotagging features.

There are several privacy implications when it comes to geotagging. One of the most dangerous is the fact that many users don't even know that they are geotagging their media when they do so; some of the applications are that transparent. For people in the company such as executives (who might carry a wealth of confidential information), this is the type of feature that should be disabled. If a potential attacker can track an executive, then the attacker can find out where the executive lives, determine when the executive is in the office, and determine the location of clients, all of which can help the attacker commit corporate espionage. When it comes down to it, the use of GPS in general should be examined carefully, weighing

the benefits against the possible vulnerabilities. This includes GPS derivatives such as GPS tagging, geofencing, and geolocation.

To be a good security administrator, you must be proactive. Preventing the threats discussed in this chapter requires updating systems and applications, and possibly redesigning networks and systems from the ground up. It also means using firewalls, host-based intrusion detection systems (HIDSs), and data loss prevention (DLP) systems. It requires in-depth configuration of applications, filtering, and secure policies. And, of course, this all signifies a need for user training.

Software is not the only place to increase security. Hardware can be physically protected, and firmware such as the BIOS should be secured as well. As mentioned, the most important thing to a company (technologically speaking) is its data. So, securing all types of storage devices, especially removable storage, is paramount. This can be done in a physical manner and in a logical manner by utilizing ***MicroSD hardware security modules (HSMs)*** and encryption, respectively.

Hardware security modules are physical devices that act as secure cryptoprocessors. This means that they are used for encryption during secure login/authentication processes, during digital signings of data, and for payment security systems. The beauty of a hardware-based encryption device such as an HSM (or a TPM) is that it is faster than software encryption.

HSMs can be found in various adapter card forms such as MicroSD HSM, as well as devices that plug into a computer via USB, and as network-attached devices. They are generally tamper-proof, giving a high level of physical security. They can also be used in high-availability clustered environments because they work independently of other computer systems and are used solely to calculate the data required for encryption keys. However, many of these devices require some kind of management software to be installed on the computer they are connected to. Some manufacturers offer this software as part of the purchase, but others do not, forcing the purchaser to build the management software themselves. Due to this lack of management software and the cost involved in general, HSMs have seen slower deployment with some organizations. This concept also holds true for hardware-based drive encryption solutions.

Often, HSMs are involved in the generation, storage, and archiving of encrypted key pairs such as the ones used in Secure Sockets Layer (SSL) sessions online, public key cryptography, and public key infrastructures (PKIs).

MDM/Unified Endpoint Management

As we discussed earlier in this chapter, MDM enables you to manage end-user mobile devices by controlling the deployment, operations, and monitoring of mobile devices used to access organization resources. It is used to enforce an organization's

security policy on mobile devices. ***Unified endpoint management (UEM)*** includes this as well as the capabilities of a mobile application management (MAM) system, mobile content management (MCM), mobile threat management (MTM), containerization, as well as identity and access management (IAM) capabilities.

Key Topic

SEAndroid

Security-Enhanced Linux (SELinux) is a security enhancement for Linux based on mandatory access control (MAC). ***SEAndroid*** is a modification of SELinux made by Google to enable the use of SELinux on the Android platform. SEAndroid provides much more robust security for the Android mobile platform by working in a deny first, least-privilege manner. This approach is, of course, great for security, but it also makes things more difficult for application developers. This is why in many cases there is an option to enable or disable SEAndroid. It is implemented in different ways by each mobile device manufacturer. For instance, on a Samsung Android device, SEAndroid utilizes a policy file to control which users and/or apps can access specific files and resources on the device. This, of course, is not something that end users should be responsible for maintaining; otherwise, they would likely end up with nonfunctioning devices. SEAndroid can be enabled in three different modes.

- Disabled mode means that SEAndroid is completely disabled on the device.
- Permissive mode runs in a manner such that the policies are in place and monitoring, but not enforcing or blocking anything.
- Enforcing mode means that SEAndroid is, in fact, blocking access to files and resources based on the policies applied.

Key Topic

Deployment Models

Around 2011, organizations began to allow employees to bring their own mobile devices into work and connect them to the organization's network (for work purposes only, of course). This bring-your-own-device concept has since grown into a more popular method of computing for many organizations. It is enticing from a budgeting standpoint but can be very difficult on you, as security administrator, and possibly on the users as well.

A computer is a computer. It doesn't matter if it's a PC from 1986 or a mobile device from this year. All computers need to be secured using the same principles and policies; however, historically mobile devices have tended to fall through the cracks. So, companies have really started gearing up the security for these devices. In most organizations, it is not feasible to stop people from bringing their smartphones into work. Some organizations have decided to embrace this practice and benefit from it with a ***bring-your-own-device (BYOD)*** policy to be used for work purposes

in addition to personal. Companies may implement similar strategies, such as ***choose-your-own-device (CYOD)***, where employees select a device from a company-approved list, or ***corporate-owned, personally enabled (COPE)***, where a company supplies employees with a phone that can also be used for personal activities.

Although these policies create a whole slew of new security considerations, some organizations are implementing BYOD and CYOD successfully by creating a well-defined demarcation point between the users' data and the organization's. When a company institutes this concept, along with an MDM solution and strong policies for theft, wireless attacks, and application security, mobile devices can survive and thrive in the enterprise or small office, yet remain safe and accessible for personal use.

NOTE Some organizations utilize a ***virtual desktop infrastructure (VDI)*** to address the BYOD challenges. This way, users can use their own hardware or corporate-owned hardware but connect to a VDI environment to access all the applications and data needed to do their work.

Secure Implementation of BYOD, CYOD, and COPE

When it comes to implementation of BYOD, CYOD, and COPE, they are all similar. For simplicity in this section, we refer to these collectively as BYOD unless CYOD or COPE needs to be addressed specifically.

First, your BYOD implementation should be driven by your policies. You should at least have a mobile device and acceptable use policy when it comes to connecting to your corporate network. The specifics of this policy should be ironed out before you begin your implementation. One of the primary concerns that should be addressed in the policy is what kind of devices you will allow on your network. Identifying these devices helps determine which is the best solution for implementing BYOD in your organization. Another determination you need to make is what kind of access you are planning to offer users of the BYOD system. For instance, if you are in a corporate environment, you may or may not allow BYOD-registered devices to access corporate data. You might decide to only allow them to access the Internet on a segmented network. If you are only allowing access to the Internet and the traffic is utilizing a segmented network that cannot interact with corporate data, your risk is much lower and you will not need to do as much validation of the device when connecting to the network. However, if the device is going to be connecting to your corporate network where it will have access to corporate data and mission-critical applications, you will want to implement your BYOD system in a much stricter manner. For example, a device connecting to a corporate network should be required to go through a multifactor authentication process. You should then

evaluate the device's security posture to determine if it meets the minimum requirements set in your policies for mobile devices. This includes minimum operating system level, up-to-date and clean antivirus/antimalware software state, validation that the device is not jailbroken, and so on.

Of course, when you are securely implementing a BYOD program, your primary goal should be to keep corporate and personal data safe. In some cases, this can and should require the devices connecting to meet some stringent requirements. However, for a BYOD program to be successful, the experience should be as easy as possible for the end user as well you, as administrator. This is where picking the right tools to implement comes in. A number of software vendors offer BYOD solutions. It is best for you to evaluate these solutions against your policies and requirements to determine which one is best for your organization. One example is the Cisco Identity Services Engine (ISE). Figure 21-1 shows the dashboard for ISE.

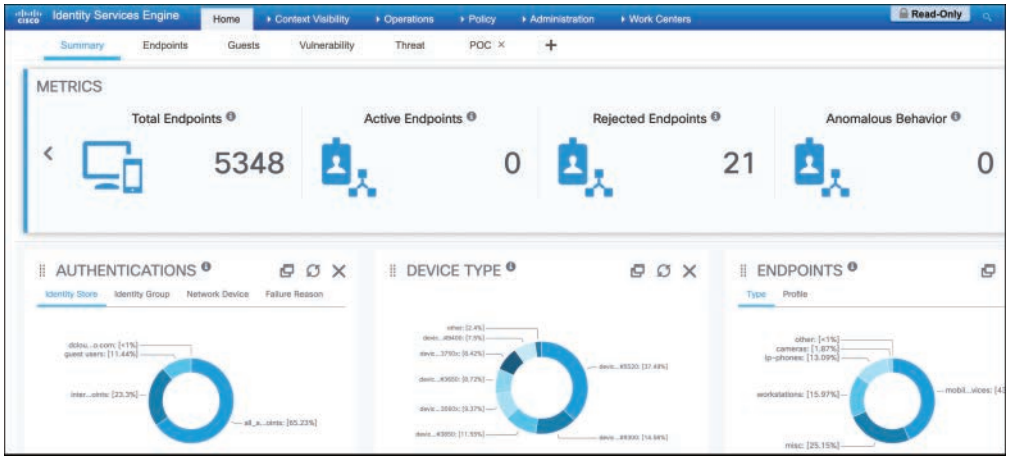


FIGURE 21-1 Cisco Identity Services Engine Dashboard View

Most BYOD platforms include a capability for implementing policy, monitoring, and onboarding devices, just to name a few. Figure 21-2 shows the BYOD work center in Cisco ISE. As you can see, it is essentially a step-by-step guide or wizard for getting started with implementing BYOD in your organization.

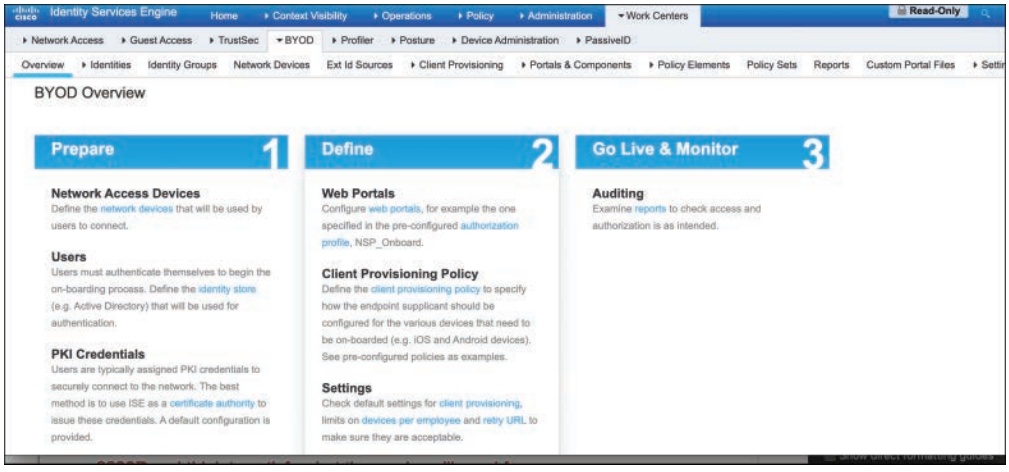


FIGURE 21-2 Cisco Identity Services Engine BYOD Work Center

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 21-3 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 21-3 Key Topics for Chapter 21

Key Topic Element	Description	Page Number
Section	Connection Methods and Receivers	570
Paragraph	Mobile device management	574
Table 21-2	Mobile Device Security Concerns and Countermeasures	580
List	Features and controls that are part of a typical MDM solution	581
Section	Mobile Devices	585
Section	SEAndroid	588
Section	Deployment Models	588

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

near field communication (NFC), radio-frequency identification (RFID), Global Positioning System (GPS), mobile device management (MDM), allow lists, block lists/deny lists, sideloading, firmware over-the-air (OTA) updates, mobile application management (MAM), MicroSD hardware security modules (HSMs), unified endpoint management (UEM), SEAndroid, bring your own device (BYOD), choose your own device (CYOD), corporate-owned, personally enabled (COPE), virtual desktop infrastructure (VDI)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What solution do some organizations use to address BYOD challenges, where users connect to an environment to access all the applications and data needed to do their work?
2. What is a security enhancement based on mandatory access control (MAC)?
3. What is the adding of data to content that would help gather location-specific information?
4. What is the denial of individual applications called?
5. What is the sending of unsolicited messages to Bluetooth-enabled devices such as mobile phones?
6. What mobile phone feature allows a phone to connect an external device such as a USB flash drive?
7. What is the unauthorized access of information from a wireless device through a Bluetooth connection?
8. What is the art of loading third-party apps from a location outside the official application store for that device?

This page intentionally left blank



This chapter covers the following topics related to Objective 3.6 (Given a scenario, apply cybersecurity solutions to the cloud) of the CompTIA Security+ SY0-601 certification exam:

- Cloud security controls
  - High availability across zones
  - Resource policies
  - Secrets management
  - Integration and auditing
  - Storage
    - Permissions
    - Encryption
    - Replication
    - High availability
  - Network
    - Virtual networks
    - Public and private subnets
    - Segmentation
    - API inspection and integration
  - Compute
    - Security groups
    - Dynamic resource allocation
    - Instance awareness
    - Virtual private cloud (VPC) endpoint
    - Container security
- Solutions
  - CASB
  - Application security
  - Next-generation secure web gateway (SWG)
  - Firewall consideration in a cloud environment
    - Cost
    - Need for segmentation
    - Open Systems Interconnection (OSI) layers
- Cloud native controls vs. third-party solutions

Applying Cybersecurity Solutions to the Cloud

In this chapter you learn how to apply cybersecurity solutions to the cloud. We start out with an overview of cloud security controls, digging deeper into the topics of high availability across zones, resource policies, and secrets management. You also gain an understanding of integration and auditing in the cloud, as well as storage permissions, encryption, replication, and high availability, then dig in to network-specific cloud security controls such as virtual networks, public and private subnets, segmentation, and API inspection and integration. Additionally, you learn the compute side of cloud security controls, including security groups, dynamic resource allocation instance awareness, virtual private cloud (VPC) endpoint, and container security. The chapter finishes by discussing solutions such as CASB, application security, next-generation secure web gateway (SWG), and firewall considerations in a cloud environment such as cost, need for segmentation, and open systems interconnection (OSI) layers. Finally, we will have a brief discussion on cloud native controls vs. third party solutions.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 22-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 22-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Cloud Security Controls	1–4
Solutions	5–7
Cloud Native Controls vs. Third-Party Solutions	8

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. In cloud computing environments, which type of policy would be used to control access to CPU and memory allocation and the like?
   a. Resource policies
   b. Key policies
   c. Wireless policies
   d. None of these answers are correct.

2. Which of the following is utilized by a cloud computing environment to handle API keys?
   a. Private key
   b. Secrets management tool
   c. Web key
   d. All of these answers are correct.

3. In cloud computing environments, what term is used for storage instances?
   a. Buckets
   b. Servers
   c. Instances
   d. Containers

4. In cloud computing environments, which type of subnet would have a route to the Internet?
   a. Private subnet
   b. Internet subnet
   c. Cloud subnet
   d. Public subnet

5. Which of the following is a tool used to control access to cloud-based environments?
   a. SCOR
   b. IAM
   c. VPC
   d. CASB

6. Which of the following solutions would help enable remote worker access more efficiently?
   a. IAM
   b. SWG
   c. CASB
   d. None of these answers are correct.

7. Which Open Systems Interconnection (OSI) layer do many cloud-based firewalls focus on?
   a. Presentation
   b. Application
   c. Network
   d. Session

8. Which type of cloud control is typically provided by the actual cloud-computing environment vendor?
   a. Third-party
   b. Cloud native
   c. Commercial
   d. Retail

Foundation Topics

Cloud Security Controls

Almost everyone these days is using the cloud or deploying hybrid solutions to host applications. The reason is that many organizations are looking to transition from capital expenditure (CapEx) to operational expenditure (OpEx). The majority of today's enterprises operate in a multicloud environment. It is obvious, therefore, that cloud computing security is more important than ever.

A multicloud environment is one in which an enterprise uses more than one cloud platform (with at least two or more public clouds), and each delivers a specific application or service. A multicloud can be composed of public, private, and edge clouds to achieve the enterprise's end goals.

NOTE Cloud computing security includes many of the same functionalities as traditional IT security. This includes protecting critical information from theft, data exfiltration, and deletion, as well as privacy.

Security Assessment in the Cloud

When performing penetration testing in the cloud, you must first understand what you can and cannot do. Most cloud service providers (CSPs) have detailed guidelines on how to perform security assessments and penetration testing in the cloud. Regardless, organizations face many potential threats when moving to a cloud model. For example, although your data is in the cloud, it must reside in a physical location somewhere. Your cloud provider should agree in writing to provide the level of security required for your customers.

Understanding the Different Cloud Security Threats

The following are questions to ask a cloud service provider before signing a contract for services:

- **Who has access?** Access control is a key concern because insider attacks are a huge risk. Anyone who has been approved to access the cloud is a potential hacker, so you want to know who has access and how they were screened. Even if these acts are not done with malice, an employee can leave and then you find out that you don't have the password, or the cloud service gets canceled because maybe the bill didn't get paid.

- **What are your regulatory requirements?** Organizations operating in the United States, Canada, or the European Union (EU) have many regulatory requirements that they must abide by (for example, ISO/IEC 27002, EU-U.S. Privacy Shield Framework, ITIL, FedRAMP, and COBIT). You must ensure that your cloud provider can meet these requirements and is willing to undergo certification, accreditation, and review.

NOTE Federal Risk and Authorization Management Program (FedRAMP) is a United States government program and certification that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMP is mandatory for U.S. federal agency cloud deployments and service models at the low-, moderate-, and high-risk impact levels. Cloud offerings such as Cisco Webex, Duo Security, Cisco Cloudlock, and others are FedRAMP certified. You can obtain additional information about FedRAMP from www.fedramp.gov and www.cisco.com/c/en/us/solutions/industries/government/federal-government-solutions/fedramp.html.

- **Do you have the right to audit?** This particular item is no small matter in that the cloud provider should agree in writing to the terms of the audit. With cloud computing, maintaining compliance could become more difficult to achieve and even harder to demonstrate to auditors and assessors. Of the many regulations touching on information technology, few were written with cloud computing in mind. Auditors and assessors might not be familiar with cloud computing generally or with a given cloud service in particular.

NOTE Division of compliance responsibilities between cloud provider and cloud customer must be determined before any contracts are signed or service is started.

- **What type of training does the provider offer its employees?** This is a rather important item to consider because people will always be the weakest link in security. Knowing how your provider trains its employees is an important item to review.
- **What type of data classification system does the provider use?** Questions you should be concerned with here include what data classification standard is being used and whether the provider even uses data classification.
- **How is your data separated from other users' data?** Is the data on a shared server or a dedicated system? A dedicated server means that your information is the only thing on the server. With a shared server, the amount of disk space, processing power, bandwidth, and so on is limited because others are sharing

this device. If it is shared, the data could potentially become comingled in some way.

- **Is encryption being used?** Encryption should be discussed. Is it being used while the data is at rest and in transit/motion? You will also want to know what type of encryption is being used. For example, there are big technical differences between DES (no longer considered secure) and AES (secure). For both of these algorithms, however, the basic questions are the same: Who maintains control of the encryption keys? Is the data encrypted at rest in the cloud? Is the data encrypted in transit/motion, or is it encrypted at rest and in transit/motion?
- **What are the service-level agreement (SLA) terms?** The SLA serves as a contracted level of guaranteed service between the cloud provider and the customer that specifies what level of services will be provided.
- **What is the long-term viability of the provider?** How long has the cloud provider been in business, and what is its track record? If it goes out of business, what happens to your data? Will your data be returned and, if so, in what format?
- **Will the provider assume liability in the case of a breach?** If a security incident occurs, what support will you receive from the cloud provider? While many providers promote their services as being unhackable, cloud-based services are an attractive target to hackers.
- **What is the disaster recovery/business continuity plan (DR/BCP)?** Although you might not know the physical location of your services, it is physically located somewhere. All physical locations face threats such as fire, storms, natural disasters, and loss of power. In case of any of these events, how will the cloud provider respond, and what guarantee of continued services is it promising?

Even when you end a contract, you must ask what happens to the information after your contract with the cloud service provider ends.

NOTE Insufficient due diligence is one of the biggest issues when moving to the cloud. Security professionals must verify that issues such as encryption, compliance, incident response, and so forth are all worked out before a contract is signed.

Cloud Computing Attacks

Because cloud-based services are accessible via the Internet, they are open to any number of attacks. As more companies move to cloud computing, look for hackers to follow. Some of the potential attack vectors that criminals might attempt include the following:

- **Denial of Service (DoS):** DoS and distributed denial-of-service attacks (DDoS) are still a threat nowadays.
- **Distributed denial-of-service (DDoS) attacks:** Some security professionals have argued that the cloud is more vulnerable to DDoS attacks because it is shared by many users and organizations, which also makes any DDoS attack much more damaging than a DoS attack.
- **Session hijacking:** This attack occurs when the attacker can sniff traffic and intercept traffic to take over a legitimate connection to a cloud service.
- **DNS attacks:** These attacks are made against the DNS infrastructure and include DNS poisoning attacks and DNS zone transfer attacks.
- **Cross-site scripting (XSS):** This input validation attack has been used by adversaries to steal user cookies that can be exploited to gain access as an authenticated user to a cloud-based service. Attackers also have used these vulnerabilities to redirect users to malicious sites.
- **Shared technology and multitenancy concerns:** Cloud providers typically support a large number of tenants (their customers) by leveraging a common and shared underlying infrastructure. This requires a specific level of diligence with configuration management, patching, and auditing (especially with technologies such as virtual machine hypervisors, container management, and orchestration).
- **Hypervisor attacks:** If the hypervisor is compromised, all hosted virtual machines could also potentially be compromised. This type of attack could also compromise systems and likely multiple cloud consumers (tenants).
- **Virtual machine (VM) attacks:** Virtual machines are susceptible to several of the same traditional security attacks as a physical server. However, if a virtual machine is susceptible to a VM escape attack, this raises the possibility of attacks across the virtual machines. A VM escape attack occurs when the attacker can manipulate the guest-level VM to attack its underlying hypervisor, other VMs, and/or the physical host.

- **Cross-site request forgery (CSRF):** This attack is in the category of web-application vulnerability and related attacks that have also been used to steal cookies and for user redirection. CSRF leverages the trust that the application has in the user. For instance, if an attacker can leverage this type of vulnerability to manipulate an administrator or a privileged user, this attack could be more severe than an XSS attack.
- **SQL injection:** This attack exploits vulnerable cloud-based applications that allow attackers to pass SQL commands to a database for execution.
- **Session riding:** Many organizations use this term to describe a cross-site request forgery attack. Attackers use this technique to transmit unauthorized commands by riding an active session using an email or malicious link to trick users while they are currently logged in to a cloud service.
- **On-path (formerly known as man-in-the-middle) cryptographic attacks:** This attack is carried out when attackers place themselves in the communication path between two users. Any time attackers can do this, there is the possibility that they can intercept and modify communications.
- **Side-channel attacks:** Attackers could attempt to compromise the cloud system by placing a malicious virtual machine in close proximity to a target cloud server and then launching a side-channel attack.
- **Authentication attacks (insufficient identity, credentials, and access management):** Authentication is a weak point in hosted and virtual services and is frequently targeted. There are many ways to authenticate users, such as factors based on what a person knows, has, or is. The mechanisms used to secure the authentication process and the method of authentication used are frequent targets of attackers.
- **API attacks:** Often application programming interfaces (APIs) are configured insecurely. An attacker can take advantage of API misconfigurations to modify, delete, or append data in applications or systems in cloud environments.
- **Known exploits leveraging vulnerabilities against infrastructure components:** As you already know, no software or hardware is immune to vulnerabilities. Attackers can leverage known vulnerabilities against virtualization environments, Kubernetes, containers, authentication methods, and so on.

TIP The Cloud Security Alliance (CSA) has a working group tasked to define the top cloud security threats. You can find details at https://cloudsecurityalliance.org/research/working-groups/top-threats. The CSA's Top Threats Deep Dive whitepaper is posted at the following GitHub Repository: https://github.com/The-Art-of-Hacking/h4cker/blob/master/SCOR/top-threats-to-cloud-computing-deep-dive.pdf. You also can find additional best practices and cloud security research articles at https://cloudsecurityalliance.org/research/artifacts/.

Key Topic

High Availability Across Zones

In cloud computing environments, high availability is addressed using the concept of regions or zones. Deploying the components of a high-availability environment across multiple zones greatly reduces the risk of an outage. Each cloud service vendor has its own way of achieving ***high availability across zones***. For instance, Amazon AWS uses elastic load balancing, whereas Microsoft Azure uses the concept of availability zones that are configured in different locations within the Azure infrastructure. There is also the concept of applying high availability across regions in each of the various cloud infrastructures.

NOTE Availability zones are isolated locations in which cloud resources can be deployed. Although resources aren't replicated across regions, they are replicated across availability zones. This setup provides load-balancing capabilities as well as high availability. Regions are separate geographic locations where infrastructure as a service (IaaS) cloud service providers maintain their infrastructure. Within regions are availability zones.

Key Topic

Resource Policies

Resource policies in cloud computing environments are meant to control access to a set of resources. A policy is deployed to manage the access to the resource itself. An example of this in Google Cloud Service is the use of constraints. These constraints are essentially a set of restrictions that are meant to outline the actions that can be controlled within the cloud environment. Often they are defined in the form of an organization policy. These policies are associated with an identity or a resource. A policy itself is what defines the permissions that are applied to the identity or resource. An identity-based policy would be utilized for a user, group, or role. This policy would define what that specific user has the permissions to do. Resource-based policies are, of course, as the name explains, associated with a specific resource available in the environment. The resource policy itself would define who has access to these resources and what they can do with them.

Key Topic

Integration and Auditing

The process of cloud integration essentially means "moving to the cloud." It is the migration of computing services from a traditional on-premises deployment to a cloud services deployment. Of course, there are many different approaches to cloud integration. Part of this integration process should include an auditing phase. This is where the organization that is "moving to the cloud" would evaluate the impact that this move will have on its current security controls as well as the privacy of data being integrated into the cloud. Each of the concerns that is highlighted from this audit phase should be addressed before moving forward with the integration process.

Key Topic

Secrets Management

In cloud computing environments, the management of things such as application programming interface (API) keys, passwords, and certificates is typically handled by some kind of ***secrets management*** tool. It provides a mechanism for managing the access and auditing the secrets used in the cloud environment. Each cloud service environment has its version of a secrets manager. In fact, the one used in Google Cloud Service is called Secret Manager. You can find additional information about the specific features of Google's Secret Manager at https://cloud.google.com/secret-manager. To access the database in Figure 22-1, the custom application will make a call to the secrets manager for the named secret. The secrets manager can manage the task of rotating the secrets and programmatically managing retrieval based on API calls. The secrets management system decrypts the secret and securely transmits the credentials to the custom application (not the database) using transport layer encryption. The application uses the credential received to access the database.

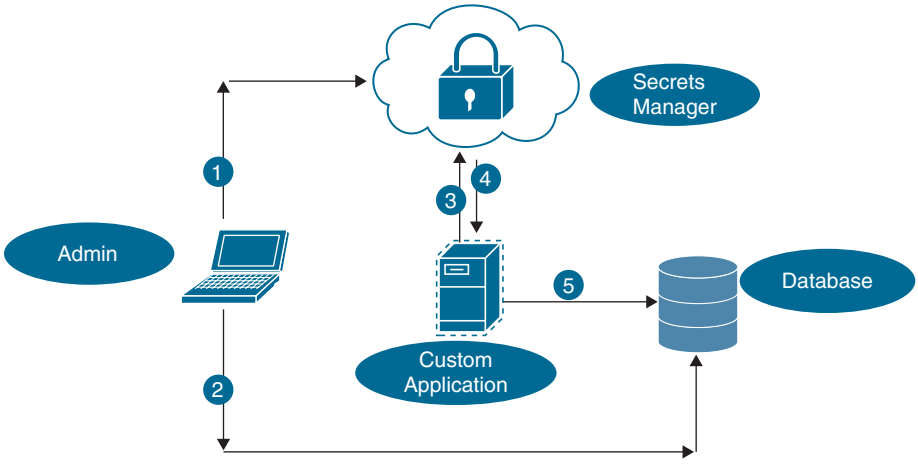


FIGURE 22-1 A Secrets Manager

Storage

The following sections cover secure storage in a cloud computing environment and how it is implemented. We provide an overview of how permissions are handled, as well as the importance of encryption, replication, and high availability when it comes to cloud storage security.

Permissions

In a cloud computing environment, storage is referred to as *buckets*. The access to these buckets is controlled by an Identity and Access Management (IAM) policy. Storage is offered as a service on the various cloud storage platforms. This service gives customers the capability to manage things like permissions and user access. Typically, the IAM service has predefined roles configured that can be applied to the resources. An example of a storage role in Google Cloud storage is the *Storage Object Creator*. This role would provide the user with permissions to create objects; however, it does not allow the user to view, delete, or replace the objects. Another example is the *Storage Object Admin* role. This role provides the user it is assigned to with full control over the objects.

Encryption

Cloud storage buckets provide various options for encryption of the data stored. Typically, cloud computing environments encrypt data stored by default; however, it is a good idea to verify this and explore the options such as user-provided keys. Before data is transferred to the cloud, it goes through a process of encoding and transforming the data. This is the process of cloud encryption. On the cloud storage provider side, the process encrypts the data, and the encryption keys are passed to the users for use in decryption. A sample method of storage using multiple layers of keys is Google's envelope encryption technique. This process involves encrypting a key with another key.

Replication

Some cloud computing environments provide a mechanism to replicate stored data between regions/zones. This provides for higher levels of redundancy and lower risk of outage or data loss. This increases the availability that we have discussed throughout the book and ties directly back to the CIA triad mentioned early on. In some cases you are actually able to specify which geographic locations you would like your data stored and replicated to. Each cloud service vendor has a different offering when it comes to options for replication of data. The main point to keep in mind when developing a replication strategy is how it will increase the durability of the data you are replicating and your overall availability.

High Availability

In cloud computing environments, the level of high availability is based on the storage class you choose and pay for. Each vendor has different specifications for the classes of storage it offers. This should be determined based on an analysis of the risk while in the integration process. This is an area where you do not want to compromise. For instance, hosting a mission-critical application that is supported by a back-end database would warrant paying for a higher level of storage availability. On the other hand, you might just be hosting some user data that is not as critical to your overall business continuity. This is an area where you might allow for a lower level of availability controls.

Network

Cloud computing environments have many options for network configuration. Each vendor has various options for segmentation and security of networks. This section begins with a look at virtual networks as well as public and private subnets. The section concludes with an overview of segmentation and API inspection and integration.

Virtual Networks

The use of the various virtual network solutions depends on how you intend to segment and/or connect the various networks. There are many use cases for virtual networks. They could be used to connect virtual machines that are dispersed between physical locations. They could also be used to set aside a segmented network for the purposes of containing specific traffic inside an enclave.

Public and Private Subnets

In cloud computing environments, the concept of a public subnet is one that has a route to the Internet. A private subnet would be a subnet in the cloud environment that does not have a route to the Internet. Obviously, a private subnet that cannot reach the Internet is limited in that way; however, many times this is the actual use case that is intended—containing traffic on a specific virtual network. On the other hand, a public subnet that is able to connect to the Internet provides additional capabilities; however, it also has a much higher risk based on the additional attack surface and exposure. A private subnet will typically be used in a situation where there is no need for the devices on that subnet to access the Internet. This makes it possible to minimize the risk of compromise by isolating these network devices. For instance, say you have a public-facing application that utilizes a back-end database. The front-end application server needs access to the Internet, but the back-end database server does not. So it's best to segment it off to a private subnet.

Segmentation

In cloud computing environments, segmentation can mean many things. There are also many ways to go about segmenting in these environments, based on the particular environment that you're working with. Some examples are the use of virtual private clouds (VPCs), which are discussed shortly, and microsegmentation of applications. As mentioned, segmentation can be used to protect devices on the network that do not need access to other parts of the network, such as the application server and back-end database previously discussed. This is an example of segmentation and how it can be implemented.

API Inspection and Integration

All cloud computing environments are able to provide ***API inspection and integration***. This allows for better automation of workflow deployment. These integrations typically need to be enabled in the environment to utilize them.

Compute

The following sections cover the concepts of compute. Some of the related concepts include security groups as well as dynamic resource allocation. We also cover instance awareness and finish up with an overview of virtual private cloud endpoint and container security.

Security Groups

Cloud-based security groups provide a mechanism for specifying separate areas where different security controls can be applied. This is similar to the use of security zones on a firewall. Security groups can be used to provide defense-in-depth to cloud-based environments. The primary function of a security group is to control the inbound and outbound traffic on the various interfaces. Typically, these interfaces would be public or private. Obviously, on the public side, it is most important to have a more strict policy for allowing inbound traffic. Most cloud providers have prebuilt security groups that can be applied. One example would be a security group that allows SSH access on port 22/TCP. Often, this is required for management purposes.

Dynamic Resource Allocation

What makes a cloud computing environment efficient is the way it is able to allocate resources based on demand. Without this capability, a cloud-based computing environment would not be feasible. ***Dynamic resource allocation*** is the concept of providing resources to tenants based on the demand. This approach provides customers

with the ability to scale resources up and down more efficiently. This allows for better resource utilization. One downside to this, however, is that there can be latency in the dynamic allocation process. This is typically not a problem unless the specific application or platform that is running on these resources is sensitive to dynamic resource allocation latency. It is best to consult the application or platform vendor to verify that this is supported before deploying in a cloud environment. A failure due to unsupported cloud deployment could result in increased downtime, which affects the availability of your systems.

Instance Awareness

Instance awareness is a concept used by cloud access security broker (CASB) solutions to enforce policies on specific parts of an application or instance. To properly enforce policy, a CASB needs to recognize the differentiation of cloud service instances—for instance, the difference between a corporate use of Dropbox and home user Dropbox access.

Virtual Private Cloud Endpoint

A ***virtual private cloud endpoint*** is a virtual endpoint device that is able to access VPC services in your cloud computing environment. Various types of VPC endpoint devices can be utilized in your environment based on the vendor. Utilizing a VPC endpoint enables you to connect to other supported services without the requirement of using a virtual private network (VPN) or Internet gateway. Consider the following best practices when using a VPC:

- Utilize availability zones to provide for high availability.
- Utilize policies to control VPC access.
- Implement ACLs and security groups where possible.
- Enable logging of API activity as well as flow logs for traffic information.

Container Security

Various tools are available for securing containerized cloud computing environments today. Some are native to the actual cloud computing environment, and some are third-party solutions that run on those environments. In today's computing environment, it is necessary to have a good handle on ***container security***. But what does that mean? First of all, visibility of what containers are running in your environment is essential. From there, you need visibility into what is actually running inside the container and what packages that container is made up of. Additionally, you need to be able to easily identify when a package inside of a container is out of

date or vulnerable and then mitigate that vulnerability. As mentioned, each vendor has a different mechanism or manager for performing these functions. If containers are going to be used in your cloud environment (you should assume they are), then you must plan for a full audit during integration. Typically, the tools for managing container security are a part of the container orchestrator tool. For instance, Google Kubernetes Engine utilizes the Google Cloud IAM for controlling and managing access. It is also important to capture the audit logs from the orchestration engine. This is something that is done by default in Google Kubernetes Engine. Additionally, it allows you to apply network control policy for managing containers and pods. These are just some of the considerations you should keep in mind when utilizing containers in your cloud environment.

Summary of Cloud Security Controls

Table 22-2 provides a summary of the key cloud security controls covered in the preceding sections.

Key Topic

Table 22-2 Key Cloud Security Controls

Cloud Security Control	Description
High availability across zones	In cloud computing environments, high availability is addressed using the concept of regions or zones. Deploying the components of a high availability environment across multiple zones greatly reduces the risk of an outage.
Resource policies	Resource policies in cloud computing environments are meant to control access to a set of resources. A policy is deployed that will manage the access to the resource itself.
Integration and auditing	The process of cloud integration essentially means "moving to the cloud." It is the migration of computing services from a traditional on-premises deployment to a cloud services deployment. Part of this integration process should include an auditing phase. This is where the organization that is "moving to the cloud" would evaluate the impact that this move will have on its current security controls as well as the privacy of data being integrated into the cloud.
Secrets management	In cloud computing environments, the management of API keys, passwords, and certificates is typically handled by some kind of secrets management tool. It provides a mechanism for managing the access and auditing the secrets used in the cloud environment.

Cloud Security Control	Description
Storage	**Permissions:** In a cloud computing environment, storage is referred to as buckets. The access to these buckets is controlled by an Identity and Access Management (IAM) policy. **Encryption:** Cloud storage buckets provide various options for encryption of the data stored. Typically, cloud computing environments encrypt data stored by default. However, it is a good idea to verify this and explore the options such as user-provided keys. **Replication:** Some cloud computing environments provide a mechanism to replicate your stored data between regions/zones. This provides for higher levels of redundancy and lower risk of outage or data loss. In some cases, you are actually able to specify which geographic locations you would like your data stored and replicated to. **High availability:** In cloud computing environments, the level of high availability is based on the storage class you choose and pay for. Each vendor has different specifications for the classes of storage it offers.
Network	**Virtual networks:** In cloud computing environments, there are many options for network configuration. Each vendor has various options for segmentation and security of networks. **Public and private subnets:** In cloud computing environments, the concept of a public subnet is one that has a route to the Internet. A private subnet would be a subnet in the cloud environment that does not have a route to the Internet. **Segmentation:** In cloud computing environments, segmentation can mean many things. There are also many ways to go about segmenting in these environments, based on the particular environment that you are working with. Some examples are the use of virtual private clouds and microsegmentation of applications. **API inspection and integration:** Most cloud computing environments are able to provide API inspection and integration. This allows for better automation of workflow deployment. These integrations typically need to be enabled in the environment to utilize them.

Cloud Security Control	Description
Compute	**Security groups:** Cloud-based security groups provide a mechanism for specifying separate areas where different security controls can be applied, similar to the use of security zones on a firewall. Security groups can be used to provide defense-in-depth to cloud-based environments. **Dynamic resource allocation:** What makes a cloud computing environment efficient is the way it is able to allocate resources based on demand. Without this ability, a cloud-based computing environment would not be feasible. Dynamic resource allocation is the concept of providing resources to tenants based on the demand. **Instance awareness:** Instance awareness is a concept used by cloud access security broker solutions to enforce policies on specific parts of an application or instance. **Virtual private cloud endpoint:** A VPC endpoint is a virtual endpoint device that is able to access VPC services in your cloud computing environment. Various types of VPC endpoint devices can be utilized in your environment based on the vendor. **Container security:** Various tools are available for securing containerized cloud computing environments today. Some are native to the actual cloud computing environment, and some are third-party solutions and run on those environments.

Solutions

The following sections cover cloud security solutions, including CASB, application security, and next-generation Secure Web Gateway (SWG). Additionally, we dive into firewall considerations in a cloud environment, such as cost, need for segmentation, and Open Systems Interconnection (OSI) layers.

CASB

A ***cloud access security broker (CASB)*** is a tool that organizations utilize to control access to and use of cloud-based computing environments. Many cloud-based tools are available for corporate and personal use. The flexible access nature of these tools makes them a threat to data leak prevention and the like. It is very easy to make the mistake of copying a file that contains sensitive data into the wrong folder and making it available to the world. This is the type of scenario that CASB solutions

help to mitigate. Of course, a number of CASB vendors are out there. Each has its own way of doing things. When securing your environment, you should first evaluate what your use cases will be for a CASB. For instance, how many cloud applications are you utilizing in your corporate environment? Make a list of each, how they are used, and by which part of the organization. You most likely will find that there are many more than you expected. When you have that information available, you can continue your journey to find a CASB vendor that will work best for your environment.

NOTE Cloud access security brokers help organizations extend on-premises security solutions to the cloud. They are solutions that act as intermediary between users and cloud service providers.

Application Security

The move to cloud-based applications has or is happening in almost every organization these days. It, of course, provides many advantages such as flexibility, ease of access, and more efficient use of resources. The downside is that it also poses some significant security concerns for these organizations. The traditional way of managing applications no longer applies. SecOps organizations need to find a way to allow access to users yet provide them with the security that they would expect from an on-prem application. The term *on-prem* (or *on-premises*) refers to any application or service that is hosted in the organization's physical facilities, not in the cloud. For some organizations, that is a strict requirement. For most, they are looking for ways to implement proper application security. This is where solutions like cloud access security broker come into play.

Aside from utilizing solutions like CASB, companies should follow additional best practices when it comes to application security. First is visibility. You need to know what applications are being accessed and used in your environment by your user base. You can check by simply monitoring your user Internet traffic to see which cloud-based applications they are using. You then need to make known and implement policies regarding what applications are allowed by your organization and how best practices on handling sensitive data can be applied. Additionally, traditional security controls such as malware protection, intrusion prevention, configuration management, and so on should be extended to your cloud-based environment. Many organizations are required to follow specific compliance when it comes to the security of their computing environment. This also applies to any cloud-based environment that they are using and should not be forgotten about. Any time you are implementing a security control on-premises, an overarching best practice is to always ask yourself "How are we doing this in the cloud?"

Next-Generation Secure Web Gateway

The concept of a next-generation ***Secure Web Gateway (SWG)*** is top of mind these days. At the time of this writing, we are currently in a global pandemic that has forced millions to work from home. For many years, the solution for providing employees a way to work from home was simply a remote-access VPN back into the office. This solution allowed all of the traffic from the employees' computers to flow back through the corporate network, which would in turn traverse the same security controls that were in place if the employees were plugged into the corporate network. With the advent of cloud-based applications and storage, this solution is no longer the most efficient way of securing remote workers. This is where the SWG comes into play. A Secure Web Gateway enables you to secure your remote workers' Internet access while not overloading the corporate Internet pipe. This approach is sometimes thought of as a cloud firewall. However, an SWG typically has many other protection mechanisms in place, including things like a CASB. One example of an SWG is Cisco Umbrella.

Firewall Considerations in a Cloud Environment

The following sections provide an overview of firewall considerations in a cloud environment. The topics we cover are cost, need for segmentation, and OSI layers.

Cost

The cost of implementing a cloud-based firewall is typically based on the amount of data processed. So, as you can imagine, this solution can become very expensive if not used efficiently. The first thing to determine is what kind of throughput you need. That starts with looking at your existing environment to determine what is actually being used. From there, you should also look at what kind of traffic makes up that bandwidth requirement. Many cloud-based firewall solutions allow you to send specific traffic through the firewall, whereas all other traffic goes directly to the Internet or takes an alternate route. This can help with reducing cost. Of course, the trade-off for this approach is that your implementation will be a bit more complex and might require more upfront cost. In the long run, however, it should save you money when it comes to cloud resources used.

Need for Segmentation

Segmentation in a cloud environment is very important for the purposes of high availability, load balancing, and disaster recovery as well as containment of security breaches. By segmenting your cloud environment, you are able to better contain a breach and more quickly mitigate it. As mentioned in previous sections, there are

many use cases for network segmentation. Understanding what is on your network is the first step in determining the need for segmentation. This requires visibility into physical and virtual devices as well as containers. This solution can be in the form of network segmentation, or it can also require microsegmentation.

Open Systems Interconnection Layers

Cloud-based firewalls can work at all layers of the OSI model; however, in many cloud computing environments, the firewall is used at the application layer to control access and mitigate threats to applications being hosted by the cloud environment. For instance, with web applications being hosted in a cloud environment, there is a need for a Layer 7 firewall, otherwise known as a web application firewall (WAF). This type of firewall is able to identify and mitigate specific web application type attacks. A WAF can also be implemented to protect a container or as a container to protect a web application in the container environment. As you can see, there are many use cases for a firewall at various layers in the cloud environment.

Summary of Cybersecurity Solutions to the Cloud

Table 22-3 provides a summary of the key cloud security solutions covered in the preceding sections.

Key Topic

Table 22-3 Key Cloud Security Solutions

Cloud Security Solution	Description
CASB	A cloud access security broker is a tool that is utilized in organizations to control access to and use of cloud-based computing environments. Many cloud-based tools are available for corporate and personal use. The flexible access nature of these tools makes them a threat to data leak prevention and the like. It is very easy to make the mistake of copying a file that contains sensitive data into the wrong folder and making it available to the world. This is the type of scenario that CASB solutions help to mitigate.
Application Secure Web Gateway	With the advent of cloud-based applications and storage, a VPN is no longer the most efficient way of securing remote workers. This is where the Secure Web Gateway comes into play. An SWG enables you to secure your remote workers' Internet access while not overloading the corporate Internet pipe. It is sometimes thought of as a cloud firewall. However, an SWG typically has many other protection mechanisms in place, including things like CASB.

Cloud Security Solution	Description
Firewall considerations in a cloud environment	**Cost:** The cost of implementing a cloud-based firewall is typically based on the amount of data processed. So, as you can imagine, this solution can become very expensive if not used efficiently. **Need for segmentation:** Segmentation in a cloud environment is very important for the purposes of high availability/disaster recovery as well as containment of security breaches. By segmenting your cloud environment, you are able to better contain a breach and more quickly mitigate. **Open Systems Interconnection (OSI) layers:** Cloud-based firewalls can work at all layers of the OSI model; however, in many cloud computing environments, the firewall is used at the application layer to control access and mitigate threats to applications being hosted by the cloud environment.

Cloud Native Controls vs. Third-Party Solutions

To put it simply, a cloud native control is typically provided by the actual cloud computing environment vendor. A non-cloud native control is provided by a third-party vendor. For instance, each cloud computing environment has security controls built into its platform. However, these controls might not be sufficient for all use cases. That is where third-party solutions come into play. Many companies out there today provide these third-party solutions to supplement where cloud native controls are lacking. Of course, most virtual machine–based controls can be deployed in any cloud computing environment; however, for it to be native, it must go through various integration, testing, and certification efforts. Because each cloud computing environment is built differently and on different platforms, the controls must be adapted to work as efficiently as possible with the specifications of the environment. Often the decision between utilizing cloud native versus third-party solutions comes down to the actual requirements of the use case and the availability of the solution that fits those requirements.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 22-4 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 22-4 Key Topics for Chapter 22

Key Topic Element	Description	Page Number
Section	High Availability Across Zones	603
Section	Resource Policies	603
Section	Integration and Auditing	604
Section	Secret Management	604
Table 22-2	Key Cloud Security Controls	609
Table 22-3	Key Cloud Security Solutions	614

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

high availability across zones, resource policies, secrets management, API inspection and integration, dynamic resource allocation, instance awareness, virtual private cloud (VPC) endpoint, container security, cloud access security broker (CASB), Secure Web Gateway (SWG)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. In cloud computing environments, which type of policy would be used to control access to things like CPU and memory allocation?
2. What cloud security control is utilized by a cloud computing environment to handle API keys?
3. In cloud computing environments, what is a term used for storage instances?
4. Which type of subnet would have a route to the Internet?
5. What tool is used to control access to cloud-based environments?
6. What cloud security solution would help to enable remote worker access more efficiently?
7. Which Open Systems Interconnection (OSI) layer do cloud-based firewalls focus on?
8. Which type of cloud control is typically provided by the actual cloud computing environment vendor?

This page intentionally left blank



This chapter covers the following topics related to Objective 3.7 (Given a scenario, implement identity and account management controls) of the CompTIA Security+ SY0-601 certification exam:

- Identity
  - Identity provider (IdP)
  - Attributes
  - Certificates
  - Tokens
  - SSH keys
  - Smart cards
- Account types
  - User account
  - Shared and generic accounts/credentials
  - Guest accounts
  - Service accounts
- Account policies
  - Password complexity
  - Password history
  - Password reuse
  - Network location
  - Geofencing
  - Geotagging
  - Geolocation
  - Time-based logins
  - Access policies
  - Account permissions
  - Account audits
  - Impossible travel time/risky login
  - Lockout
  - Disablement

Implementing Identity and Account Management Controls

In this chapter you learn about implementing identity and account management controls. The chapter starts by examining identity concepts such as identity provider (IdP), attributes, certificates, tokens, SSH keys, and smart cards. It continues with an overview of account types, discussing user accounts, shared and generic accounts, guest accounts, and service accounts. The chapter concludes with a discussion of account policies, including password complexity, password history, password reuse, network location, geofencing, geotagging, and geolocation. This discussion also includes topics such as time-based logins, access policies, account permissions, account audits, impossible travel time/risky login, lockout, and disablement.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 23-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 23-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Identity	1–4
Account Types	5–7
Account Policies	8–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which type of service provider also manages the authentication and authorization process on behalf of other systems?
   a. Service provider (SP)
   b. Identity provider (IdP)
   c. Principal
   d. None of these answers are correct.

2. Which of the following is best defined as the process of providing the identity of a subject or user?
   a. Identity
   b. Authentication
   c. Authorization
   d. All of these answers are correct.

3. Which standard format is used for most certificates?
   a. X.509
   b. EFS
   c. X.409
   d. PEM

4. Which of the following is an open standard for exchanging authentication and authorization data between identity providers?
   a. RSA
   b. RADIUS
   c. Duo
   d. SAML

5. Which of the following types of accounts would be used for a web server?
   a. Server account
   b. Generic account
   c. User account
   d. Service account

6. Which tool can be used to validate who is logged in to a Windows system?
   a. Winlogon
   b. PSLoggedOn
   c. whoami
   d. None of these answers are correct.

7. Which command can be used on a Linux system to determine who is logged in?
   a. **WinlogON**
   b. **PSlogOn**
   c. **whoami**
   d. **whoareyou**

8. Which method can be used in determining the physical location of a user trying to authenticate?
   a. Username
   b. Geolocation
   c. Group
   d. Domain

9. Which attribute would be used if you wanted to restrict when a user is able to log in?
   a. Geolocation
   b. Geotagging
   c. Time-based
   d. None of these answers are correct.

10. Which of the following is the process of information being attached in metadata to files such as pictures?
    a. Geolocation
    b. Geotagging
    c. Lockout
    d. None of these answers are correct.

Foundation Topics

Identity

Key Topic

Identification and authentication are often performed together; however, it is important to understand that they are two different operations. Identification is about establishing who you are, whereas authentication is about proving you are the entity you claim to be. Here, we discuss some of the concepts related to ***identity***, starting with the use of an identity provider (IdP). We also cover attributes, certificates, tokens, SSH keys, and smart cards.

Identity Provider (IdP)

A further evolution of the SSO model within a single organization is a model where a user can authenticate once and then have access to resources across multiple organizations not managed under the same IAM system. A *federation* is a collection of distinct organizations that agree to allow users to use one set of credentials for authentication and authorization purposes. The identity used by the users across organizations is called a *federated identity*.

At the base of the federation is the concept of *trust* between the organization entities. In fact, each organization should trust that the authentication and authorization process is carried out in a secure way by the party providing that service.

The concept of federation has been further formalized by introducing the following concepts:

- **Principal:** The end user who requests service from a service provider and whose identity can be authenticated.
- **Service provider (SP):** In some cases also called the relying party (RP). The SP is defined as the system entity that provides service to the principal or other entities in the federation.
- ***Identity provider (IdP)*:** The service provider that also manages the authentication and authorization process on behalf of the other systems in the federation.
- **Assertion:** The information produced by the authentication authority (for example, the IdP). It is usually provided to the SP to allow the user to access its resource. The assertion proves that the user has been authenticated and includes additional user attributes and authorization directives.

In a federation context, an SP can rely on multiple IdPs, and one IdP can serve multiple SPs. When a user wants to access resources with one SP, the SP determines

which IdP to use to authenticate the user. The choice happens based on the user identifier or preference (for example, the user may indicate a specific IdP), or the choice happens based on the domain name associated with the user email address. This process is called *discovery of identity*.

The SP then redirects the user to the IdP for the authentication process. After the user is authenticated, the IdP generates an assertion that proves the identity and includes additional user and authorization information.

In the example in Figure 23-1, the user authenticates with an SSO server that is in a different organization than the one in Server B, which provides service to the user it belongs to. In this case, the SSO server acts as the IdP, and Server B is the SP.

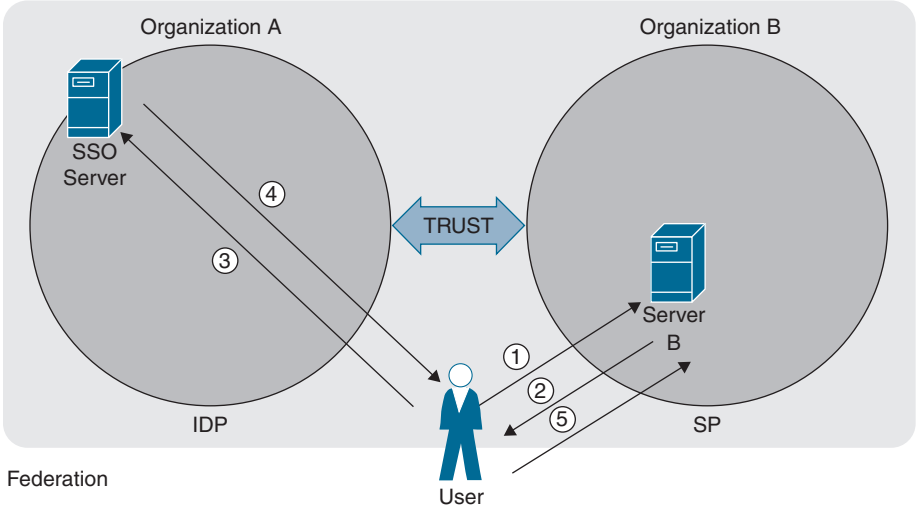


FIGURE 23-1 Federated SSO

The user sends a request to Server B (step 1), which redirects the user to the SSO server for authentication (steps 2 and 3). The user then authenticates with the SSO server and receives proof of authentication, the assertion, which is provided to Server B (steps 4 and 5). Server B, after verifying the information in the assertion, grants access to resources.

Several protocols and frameworks are currently used to implement SSO and identity federation: SAML, OAuth2, and OpenID Connect are popular examples.

Authentication

Authentication is the process of proving the identity of a subject or user. Once a subject has identified itself in the identification step, the enforcer has to validate the identity—that is, be sure that the subject (or user) is the one it is claiming to be. This is done by requesting that the subject (or user) provide something that is unique to the requestor. This could be something known only by the user, usually referred to as *authentication by knowledge*, or owned only by the user, usually referred to as *authentication by ownership*, or it could be something specific to the user, usually referred to as *authentication by characteristic*.

Authentication by Knowledge

In authentication by knowledge, the user provides a secret that is only known by him or her. An example of authentication by knowledge would be a user providing a password, giving a personal identification number (PIN) code, or answering security questions.

The disadvantage of using this method is that once the information is lost or stolen (for example, if a user's password is stolen), an attacker would be able to successfully authenticate.

Authentication by Ownership

In authentication by ownership, the user is asked to provide proof that he or she owns something specific—for example, a system might require an employee to use a badge to access a facility. Another example of authentication by ownership is the use of a ***token*** or ***smart card***. A ***certificate*** is also a factor considered to be something you have. These are used on smart cards as well as for use cases like ***SSH keys***. The sections that follow delve more deeply into these methods of providing identity.

Authentication by Characteristic Attributes

A system that uses authentication by characteristic authenticates the user based on some physical or behavioral characteristic, sometimes referred to as a ***biometric attribute***. Here are the most used physical or physiological characteristics:

- Fingerprints
- Facial recognition
- Retina and iris
- Palm and hand geometry
- Blood and vascular information
- Voice

Here are examples of behavioral characteristics:

- Signature dynamic
- Keystroke dynamic/pattern

The drawback of a system based on this type of authentication is that it's prone to accuracy errors. For example, a signature-dynamic-based system would authenticate a user by requesting that the user write his or her signature and then comparing the signature pattern to a record in the system. Given that the way a person signs his or her name differs slightly every time, the system should be designed so that the user can still authenticate even if the signature and pattern are not exactly the one in the system. However, it should also not be too loose and thus authenticate an unauthorized user attempting to mimic the pattern.

Two types of errors are associated with the accuracy of a biometric system:

- A **Type I error**, also called *false rejection rate (FRR)*, happens when the system rejects a valid user who should have been authenticated.
- A **Type II error**, also called *false acceptance rate (FAR)*, happens when the system accepts a user who should have been rejected (for example, an attacker trying to impersonate a valid user).

The *crossover error rate (CER)*, also called the *equal error rate (EER)*, is the point where FFR and FAR are equal. This is generally accepted as an indicator of the accuracy (and hence the quality) of a biometric system.

Table 23-2 lists the three authentication methods described in this section and provides a short description and examples of each.

Key Topic

Table 23-2 Authentication Methods

Authentication Method	Description	Examples
Authentication by knowledge	Something the user knows	Password, PIN
Authentication by ownership	Something the user owns	Smart card, badge, token
Authentication by characteristic	Something the user is or does	Fingerprint, hand geometry, keystroke dynamic

Certificates

An identity certificate describes the client and contains the public key of an individual host (the client). Identity certificates are used by web servers, APIs, VPN clients, and web browsers (in some cases).

For instance, a web application using the TLS protocol for encryption over HTTPS is assigned an identity certificate. The certificate you see in your browser for that

web application (or website) includes the public key for the server. Any device that wants to verify a digital signature must have the public key of the sender. Your browser verifies the certificate authority's (CA's) signature on the digital certificate of the web application also using the public key of the root certificate authority that issued the certificate. In practice, this public key for the root CA is built in to most browsers. Figure 23-2 shows how authenticating and enrolling with a CA works.

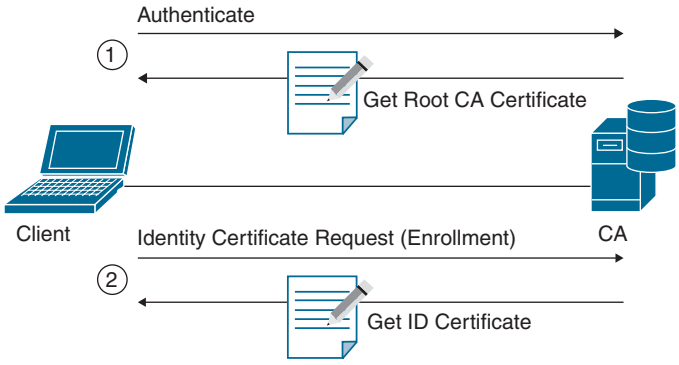


FIGURE 23-2 Authenticating and Enrolling with a CA

Tokens

A one-time password (OTP) is a randomly generated password that can be used only once. One of the most-used methods for implementing one-time password authentication is through a token device. The token device can be either a hardware device or implemented in software (soft-token), and it is basically a one-time password generator. For example, most of the authentication systems for online banking use token technologies.

A token device can work in two ways: synchronously and asynchronously. In most cases, the token generator is protected through a password or PIN. In synchronous token authentication, the token generator is synchronized with the authenticator server (for example, via time synchronization). When a user needs to authenticate, he or she uses the token to generate a one-time password that's used to authenticate to the system. In an asynchronous token system, the authenticator produces a challenge. The user inputs the challenge in the token generator, which uses that information to generate the one-time password.

NOTE Hardware-based security tokens are physical devices given to authorized users to help with authentication. These devices might be attached to a keychain or might be part of a card system. Hardware-based tokens might be used as part of the door access system or as something that gives access to an individual computer. As one example, RSA tokens carry and generate rolling one-time passwords, each of which is valid for only one login session or transaction.

SSH Keys

Secure Shell (SSH) is a protocol that can create a secure channel between two computers or network devices, enabling one computer or device to remotely control the other. Designed as a replacement for Telnet, it is commonly used on Linux and UNIX systems, and nowadays also has widespread use on Windows clients. It depends on public key cryptography to authenticate remote computers. One computer (the one to be controlled) runs the SSH daemon, while the other computer runs the SSH client and makes secure connections to the first computer (which is known as a server), as long as a certificate can be obtained and validated.

As mentioned previously, understanding the terminology is important for you in mastering encryption and secure communication technologies. Figure 23-3 explains the key components and their functions as well as provides examples of their implementation.

Symmetric Encryption Algorithms	Asymmetric Encryption Algorithms	Digital Signatures	Diffie-Hellman Key Exchange
Use the same key for encrypting and decrypting data.	Use a public and private key. One key encrypts the data and the other is used to decrypt.	Encryption of hash using private key, and decryption of hash with the sender's public key.	Uses a public-private key pair asymmetrical algorithm, but creates final shared secrets (keys) that are then used by symmetrical algorithms.
AES, DES, 3DES, IDEA	RSA, Diffie-Hellman	RSA Signatures	Used in IPsec

Confidentiality	Data Integrity	Authentication
Encryption algorithms provide this by converting clear text into cipher text.	Validates data by comparing hash values.	Verifies the peer's identity.
AES, DES, 3DES RSA, IDEA	MD5, SHA-1, SHA-2, SHA-3	PSKs, RSA Signatures

FIGURE 23-3 Key Implementation Examples

Smart Cards

The smart card falls into the category of "something a person has" and is known as a token. It's the size of a credit card and has an embedded chip that stores and transacts data for use in secure applications such as hotel guest room access, prepaid phone services, and more. Smart cards have multiple applications, one of which is to authenticate users by swiping the card against a scanner, thus securing a computer or a computer room. The smart card might have a photo ID as well. Examples of smart cards include the Personal Identity Verification (PIV) card, which is required for all U.S. government employees and contractors, and the Common Access Card (CAC), which is used to identify Department of Defense (DoD) military personnel, other DoD civilian government employees, and so on.

These cards not only identify the person and are responsible for authentication to buildings and systems but can also encrypt and digitally sign emails. These cards might be used as part of a multifactor authentication scheme in which there is a combination of username/password (or PIN) and a smart card. Advanced smart cards have specialized cryptographic hardware that can use algorithms such as RSA and 3DES but generally use private keys to encrypt data. A smart card might incorporate a microprocessor (as is the case with the PIV and CAC cards). A smart card security system usually is composed of the smart card itself, smart card readers, and a back-office database that stores all the smart card access control lists and history.

Older technologies use proximity sensors, but this is not considered very secure today. However, the more complex the technology, the more it costs. Often, in these situations, budgeting becomes more important to organizations than mitigating risk; and generally, the amount of acceptable risk increases as the budget decreases. So, you will probably see proximity-based door access systems. HID (also known as HID Global) is an example of a company that offers various levels of door access control systems.

Account Types

Identifying who is logged in to a system is important for knowing how the system will be used. Administrators typically have more access to various services than other users because their job requires those privileges. Employees within Human Resources might need more access rights than other employees to validate whether an employee is violating a policy.

Guest accounts typically require very few access rights because they are considered a security risk to most organizations. ***Shared and generic accounts*** are shared between multiple users who share the same password. This is not an ideal way to handle ***user account*** creation and can cause you to lose accounting data. A ***service account*** is typically used on a server to provide a separate set of credentials and permissions to an application or service that is running. For instance, a server that is running Apache web

server might have an "apache_user" account. That account would then be provided with only the access it needs to be able to perform the functions that service provides.

In summary, best practice for provisioning access rights is to enforce the concept of ***least privilege***, meaning to provision the absolute least amount of access rights required to perform a job. It is also very important to regularly complete account audits to verify that users do not have more access than needed. Additionally, an account audit should verify that an account is actually required for use. If it is not, then it should always be disabled.

People can be logged in to a system in two ways:

- The first method is to be physically at a keyboard logged in to the system.
- The other method is to remotely access the system using something like a Remote Desktop Connection (RDP) protocol. Sometimes the remote system is authorized and controlled, such as using a Citrix remote desktop solution to provide remote users access to the desktop, whereas other times it's a malicious user who has planted a remote access tool (RAT) to gain unauthorized access to the host system.

Identifying post-breach situations is just one of the many reasons that monitoring remote connections should be a priority for protecting your organization from cyber breaches.

You can identify who is logged in to a system through a few different approaches. For Windows machines, the first method involves using the Remote Desktop Services Manager suite. This approach requires the software to be installed. When the software is running, you can remotely access the host to verify who is logged in.

Another tool you can use to validate who is logged in to a Windows system is the PsLoggedOn application. For this application to work, it has to be downloaded and placed somewhere on your local computer that will be remotely checking hosts. After it's installed, you simply open a command prompt and execute the command **C:\PsTools\psloggedon.exe\\HOST_TO_CONNECT**.

You can use Windows PowerShell to obtain detailed information about users logged in to the system and many other statistics that can be useful for forensics and incident response activities. Similarly to the psloggedon.exe method, you can take advantage of PowerShell modules like Get-ActiveUser, which is available at www.powershellgallery.com/packages/Get-ActiveUser/1.4.

NOTE The PowerShell Gallery (www.powershellgallery.com) is the central repository for PowerShell modules developed by Microsoft and the community.

For Linux machines, various commands can show who is logged in to a system, such as the **w**, **who**, **users**, **whoami**, and **last "user name"** commands. Example 23-1 shows the output of the **w** command. Two users are logged in as **omar** and **root**.

Example 23-1 Using the Linux **w** Command

```
$w
 21:39:12 up 2 days, 19:18,  2 users,  load average: 1.06, 0.95, 0.87
USER    TTY       FROM          LOGIN@   IDLE    JCPU    PCPU WHAT
root    tty1      -             18:29    0.00s   0.02s   0.00s ssh 10.6.6.3
omar    pts/2     10.6.6.3      21:36    0.00s   0.01s   0.00s w
```

Example 23-2 shows the output of the **who**, **users**, **last**, and **lastlog** Linux commands. In Example 23-2, you can see the details on the users root and omar, as well as when they logged in to the system using the **last** command. The **lastlog** command reports the most recent login of all users or of a given user. In Example 23-2, the **lastlog** command displays information of all users in the system.

Example 23-2 The **who** and **users** Linux Commands

```
$who
root      tty1           Jun  3 18:29
omar      pts/2          Jun  3 21:36 (10.6.6.3)
$users
omar root
$last
omar      pts/2       10.6.6.3       Wed Jun  3 21:36   still logged in
root      tty1                       Wed Jun  3 18:29   still logged in
wtmp begins Wed Jun  3 18:29:09 2020
$lastlog
Username     Port     From           Latest
root         tty1                    Wed Jun  3 18:29:09 +0000 2020
daemon                               **Never logged in**
bin                                  **Never logged in**
sys                                  **Never logged in**
sync                                 **Never logged in**
games                                **Never logged in**
man                                  **Never logged in**
lp                                   **Never logged in**
mail                                 **Never logged in**
news                                 **Never logged in**
```

```
uucp                                  **Never logged in**
proxy                                 **Never logged in**
www-data                              **Never logged in**
backup                                **Never logged in**
list                                  **Never logged in**
irc                                   **Never logged in**
gnats                                 **Never logged in**
nobody                                **Never logged in**
syslog                                **Never logged in**
systemd-network                       **Never logged in**
systemd-resolve                       **Never logged in**
messagebus                            **Never logged in**
postfix                               **Never logged in**
_apt                                  **Never logged in**
sshd                                  **Never logged in**
uuidd                                 **Never logged in**
chelin       pts/2    10.6.6.45     Tue Jun  2 11:22:41 +0000 2020
omar         pts/2    10.6.6.3      Wed Jun  3 21:36:53 +0000 2020
```

Each option shows a slightly different set of information about who is currently logged in to a system. One command that displays the same information on a Windows system is the **whoami** command.

Many administrative tools can be used to remotely access hosts, so the preceding commands can be issued to validate who is logged in to the system. One such tool is a virtual network computing (VNC) server. This method requires three pieces. The first part is having a VNC server that will be used to access clients. The second part is having a VNC viewer client installed on the host to be accessed by the server. The final part is an SSH connection that is established between the server and client after things are set up successfully. SSH can also be used directly from one system to access another system using the **ssh "remote_host"** or **ssh "remote_username@remote_host"** command if SSH is set up properly. Many other applications, both open source and commercial, can provide remote desktop access service to host systems.

It is important to be aware that validating who is logged in to a host can help identify when a host is compromised. According to the kill chain concept, attackers who breach a network will look to establish a foothold through breaching one or more systems. When they have access to a system, they can seek out other systems by pivoting from system to system. In many cases, attackers want to identify a system with more access rights so that they can increase their privilege level, meaning gain access to an administration account, which typically can access critical systems. Security tools

that include the ability to monitor users logged in to systems can flag whether a system associated with an employee accesses a system that's typically accessed only by administrator-level users, thus indicating a concern for an internal attack through a compromised host. The industry calls this type of security breach *detection*, meaning technology looking for post-compromise attacks.

Be sure to keep these key concepts in mind:

- Employing least privilege means to provision the absolute minimum amount of access rights required to perform a job.
- The two methods to log in to a host are locally and remotely.
- Common methods for remotely accessing a host are using SSH and using a remote access server application such as VNC.

Key Topic

Account Policies

Next, we cover account policies and some of the important aspects that should be considered. We start by covering password complexity, password history, and password reuse. Next, we address network location, including the concepts of geofencing, geotagging, and geolocation. From there, we provide an overview of time-based logins, access policies, account permissions, and account audits. We conclude with a look at impossible travel time/risky login, lockout, and disablement.

Introduction to Identity and Access Management

Identity and access management (IAM) has a very broad definition and in general includes all policies, processes, and technologies used to manage the identity, authentication, and authorization of an organization's resources. Several disciplines and technologies are usually covered under the umbrella of IAM: access control password management, the IAM lifecycle, directory management, and single sign-on (SSO), among others.

Phases of the Identity and Access Lifecycle

One of the properties of a secure identity is the secure issuance of that identity. Additionally, access privileges should be associated with an identity, and the identity's validity and permissions should be constantly reviewed. At times, an identity and permissions should be revoked, and a process should be established to do this in a secure way. These processes are called identity proof and registration, account provisioning, access review, and access revocation. They all go under the umbrella of identity and account lifecycle management.

Figure 23-4 shows the following four phases of the identity and access lifecycle:

- **Registration and identity validation:** A user provides information and registers for a digital identity. The issuer verifies the information and securely issues a unique and nondescriptive identity.
- **Privileges provisioning:** The resource owner authorizes the access rights to a specific account, and privileges are associated with it.
- **Access review:** Access rights are constantly reviewed to avoid privilege creep.
- **Access revocation:** Access to a given resource may be revoked due, for example, to account termination.

Key Topic

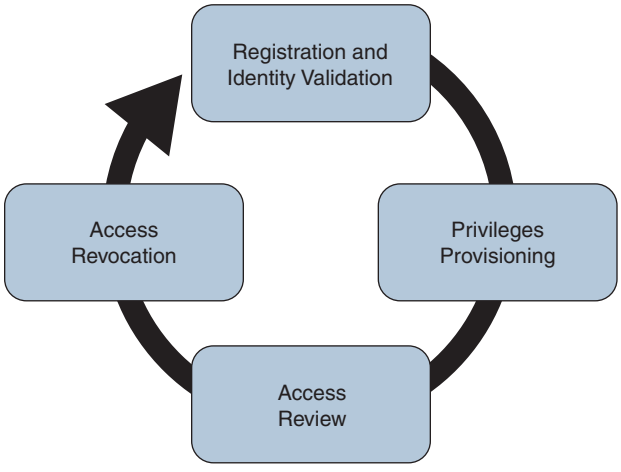


FIGURE 23-4 Identity and Access Lifecycle

Let's review each of these phases in a bit more detail.

Registration and Identity Validation

The first step in a secure identity lifecycle is user registration. During this phase, a user registers data to request an identity. The second step of this process is to verify the identity. This can be done in several ways, depending on the privileges associated with that identity. For example, starting the identity validation for a system administrator may require additional steps compared to those for a normal user. During this phase, a user could be asked to provide an identity card, HR could perform a background check during the onboarding process, proof of a specific clearance level could be requested, and so on. Finally, the identity assigned is unique and nondescriptive.

Privileges Provisioning

After an identity has been assigned, privileges or access rights should be provisioned to that account. The privileges should be assigned by using the main security principles discussed throughout this book—that is, least privileges, separation of duty, and need to know. In general, privileges are assigned in accordance with the organization's security policy.

Depending on the access control model applied, the process might need to ensure that an authorization request is sent to the resource owner and that privileges are not assigned until the access has been approved. A temporal limit should also be applied to the privileges assigned.

For highly sensitive privileges, a more formal process might need to be established. For example, users may be asked to sign a specific nondisclosure agreement. Provisioning could also apply to existing accounts requesting access to additional resources, for example, due to a job change within the organization.

NOTE The registration, identity validation, and privileges provisioning phases are grouped together under the account provisioning step.

Access Review

Access rights and privileges associated with an account should be constantly reviewed to ensure that there is no violation to the organization's security policy. The process should ensure a regular review of privileges as well as an event-driven review, such as when a user changes roles.

One of the issues in large organizations is the unneeded assignment of privileges, which brings up the *privileges creep* issue.

Access Revocation

When an employee changes jobs or leaves the organization, the organization might need to partially or completely revoke the employee's associated access rights. A formal process should be established (offboarding) to make sure this is done properly. In some cases, privileges may need to be revoked before the actual event (for example, an involuntarily job termination) to ensure the employee does not cause damage to the organization before officially leaving.



Password Management

A password is a combination of characters and numbers that should be kept secret, and it is the most common implementation of authentication by knowledge. Password authentication is usually considered one of the weakest authentication methods, yet it's one of the most used due to its implementation simplicity.

The weakness of password authentication is mainly due to the human factor rather than technological issues. Here's a list of some typical issues that lead to increased risk when using passwords as the sole authentication method:

- Users tend to use the same password across all systems and accounts.
- Users tend to write down passwords (for example, on a sticky note).
- Users tend to use simple passwords (for example, their child's name or 12345).
- Users tend to use the default system password given at system installation.

Password management includes all processes, policies, and technologies that help an organization and its users improve the security of their password-authentication systems. Password management includes policies and technologies around password creation, password storage, and password reset.

Password Creation

One of the most important steps in password management is creating a standard to define secure password requirements. This standard needs to be applied across the organization and for all systems. An organization should take into consideration the following requirements when building policies, processes, and standards around password creation:



- **Password complexity:** Establishing a policy about password strength is very important to reduce the risk of users setting up weak passwords, which are easier to compromise via brute-force attacks, for example. Complexity and length requirements contribute to increasing the strength of a password. Complexity is usually enforced by asking users to use a combination of characters, numbers, and symbols. Password length increases the difficulty of cracking a password. The shorter the password, the higher the risk. The strength and entropy of a password are the main factors used to measure the quality of a password. NIST SP 800-63b provides more information about password entropy and how passwords can be used in electronic authentication systems. Due to continued research on password use, NIST has recently updated its password guidelines. The following are some of the recent updates made to password requirements in NIST SP 800-63b.

- Users no longer have to use special characters.
- Users should be able to use all characters.
- Copying and pasting passwords is acceptable.
- Password policies should not require employees to change passwords on a regular basis.
- Password fields should allow for up to at least 64 additional characters on top of the required 8.

- **Password history:** The age of a password (or better, the maximum age of a password) is an important attribute. Changing a password frequently is considered a best practice. The longer a password is used, the higher the risk of password compromise. The password requirement policy should dictate the maximum age of a password. Changing passwords frequently is better for security; however, it creates additional administrative overhead for users and systems.
- **Password reuse:** Reusing the same password or part of it also increases the risk of password compromise. It is common practice to change just the last digit of a password or to use only two passwords repeatedly and swap them when the time comes. Policy around reusability should ensure that passwords are not reused within a given amount of time.

The policies around the creation of a password should also specify whether the password is created by the user or is automatically generated by the system. A hybrid approach would use both methods by combining a user-chosen password with a system-generated one. Table 23-3 summarizes the pros and cons of each of these methods.

Key Topic

Table 23-3 Password-Generation Methods

Method	Description	Pros	Cons
User-generated password	Users generate the password themselves.	Simple to remember.	Usually leads to an easily guessable password. Users may reuse the same password on multiple systems.
System-generated password	The password is generated by the system.	Strong password. Compliant with security policy.	Difficult to remember. Users tend to write down the password, thus defeating the purpose.

Method	Description	Pros	Cons
OTP and token	The password is generated by an external entity (such as hardware or software) that is synchronized with internal resources. The device is usually protected by a user-generated password.	Users do not need to remember a difficult password.	More complicated infrastructure. It makes use of hardware or software to generate the token, which increases maintenance and deployment costs.

User-Generated Password

Using passwords created by the user is the easiest method but is the riskiest from a security point of view. Users tend to use easy passwords, reuse the same passwords, and, in some cases, disclose passwords to others. Enforcing password requirements helps reduce the risk.

System-Generated Password

Using system-generated passwords is a stronger method than using user-created passwords because the password requirements are directly enforced. In most cases, the system can create the passwords by using a random password generator, which ensures higher entropy and is usually more difficult to compromise. The drawback of this method is that these types of passwords are more difficult to remember. Users, therefore, tend to write them down, which defeats the purpose of having a secure password.

Attribute-Based Access Control (ABAC)

Attribute-based access control (ABAC) is an access model that is dynamic and context-aware. Access rights are granted to users through the use of multiple policies that can combine various user, group, and resource attributes together. This type of control makes use of IF-THEN statements based on the user and requested resource. For example, *if* David is a system administrator, *then* allow full control access to the \\dataserver\adminfolder share. If this solution is implemented properly, it can be more flexible. As of this writing, many technologies—and organizations—are moving toward a more context-sensitive, context-aware mindset when it comes to authentication and access control.

Table 23-4 summarizes some of the attributes and identity terms that are used in account policies.

Key Topic

Table 23-4 Attributes and Identity Terms Used in Account Policies

Attribute/Term	Key Points
Network location	The location where users are connected to the network can be used to determine what access they have. For instance, a doctor connected to the guest network in the lobby should not have access to medical records. This is how the concept of ***geofencing*** works in identification and authorization.
Geolocation	Geolocation is the actual method of determining the physical location of the user trying to authenticate. This is often used in the implementation of mobile device policies, where users are only allowed to access a network based on the location they are reaching the Internet from.
Geotagging	Geotagging is the process of attaching location information in the metadata of files such as pictures taken from a smartphone. The implementation of a policy using geotagging would typically be to restrict the use of geotagging for sensitive locations or data.
Time-based logins	Time is another attribute that can be used when authenticating and authorizing users. When users log in and provide their identity, they can be given specific access based on the time they connect. For instance, if you do not expect that someone should be connecting to your wireless network at 3 a.m., then you can set a policy that blocks access between specific hours of operation.
Disablement/lockout	An account can be disabled or locked out based on specific attributes—for instance, if a user is known to be located in the U.S. but connects from China. This may be an indication that the user credentials have been compromised and should be temporarily disabled until this access can be verified.
Account audits	This is the process of auditing the permissions that are assigned to accounts on a system. Many times, permissions are added when needed but never removed when not needed.
Impossible travel time/risky login	This is a concept related to geolocation and geofencing where a potential compromise is identified based on the fact that it would be impossible for a user to be in two places at once—for instance, if you authenticated at 3 p.m. EST from New York, then 10 minutes later tried to authenticate from Florida. This access would not be possible and is an indication of account compromise.

NOTE Another type of access control method is known as anonymous access control—for example, access to an FTP server. This method uses attributes before access is granted to an object. Authentication is usually not required.

NOTE In general, access control can be centralized or decentralized. *Centralized* access control means that one entity is responsible for administering access to resources. *Decentralized* access control means that more than one entity is responsible, and those entities are closer to the actual resources than the entity would be in a centralized access control scenario.

Rights, Permissions, and Policies

Now that you have a plan for access control, you need to implement it in a tangible way. By strategically setting up organizational units, users, and groups, and by assigning permissions according to a chosen access control model, you can create a safe, guarded working area for all employees. In so doing, you can protect the data on the network.

Users, Groups, and Account Permissions

User accounts can be added to individual computers or to networks. For example, a Windows client, Linux computer, or macOS system can have multiple users. And larger networks that have a controlling server, for example, a Windows domain controller, enable user accounts that can access one or more computers on the domain. In a Microsoft domain, users are added in Active Directory Users and Computers (ADUC), as shown in Figure 23-5.

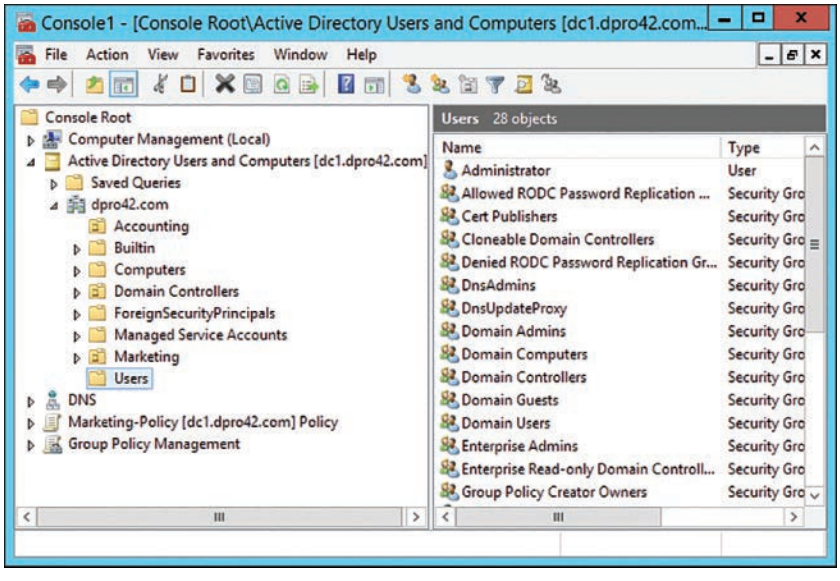


FIGURE 23-5 The Users Folder Within ADUC on a Windows Server

ADUC can be accessed from Administrative Tools or added as a snap-in to a Microsoft Management Console (MMC). Users can be added in one of two places:

- **In the Users folder:** This folder is located inside the domain name within ADUC.
- **In an OU:** Organizational units can be created within the domain. They are often made to mimic the departments of a company. In Figure 23-5, there are Accounting and Marketing OUs; users can be created within these OUs.

User rights can be modified within the particular user's Properties window. Many more rights are associated with a user account that is stored on a Windows Server domain controller than are on an individual Windows client computer. For example, the Account tab can be configured so that the user account has an expiration date. You can see this in Figure 23-6, where at the bottom of the Properties window, Alice's account has been configured to expire on April 13, 2022. Immediately after that expiration date, the user cannot log on to the domain unless her account is reconfigured or she logs on as someone else.

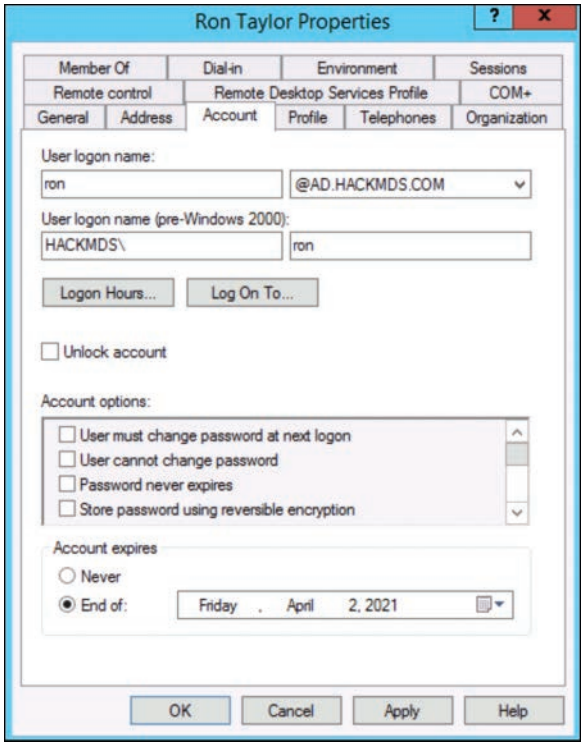


FIGURE 23-6 User Account Expiration Date

NOTE Users cannot log on to a network after their account has expired. The Account Expiration date in Windows controls this access.

By clicking the Logon Hours button, you can configure time-of-day restrictions so that a user can log on only at certain times throughout the week. An example is shown in Figure 23-7 where Alice's user account has been configured in such a way that she can log on to the domain only between 8 a.m. and 6 p.m. Monday through Friday. If she attempts to log on at any other time, the system will deny access. These kinds of access rights are available on domain controllers.

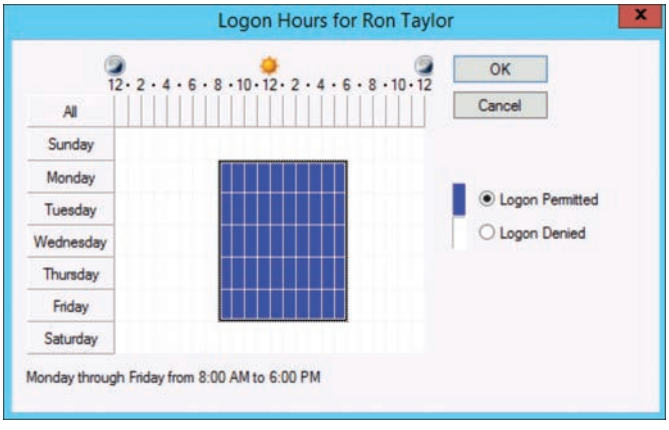


FIGURE 23-7 Time-of-Day Policies for a Standard User

NOTE Users can log on to the network only during their configured logon hours.

Sometimes users have more than one account. This might have been done to allow access to multiple systems or resources. There are plenty of different issues that can occur because of this setup. To mitigate problems that can develop from a user having two accounts, consider the consolidation of accounts, for example, utilizing a federated identity management (FIM) system, one that will incorporate single sign-on. User administration can also benefit from credential management, where passwords, certificates, and other logon credentials are stored in a special folder called a vault. A security administrator should also consider the use of roles (RBAC) and user groups.

Groups can be created to classify users and to ease administration when assigning permissions. In Figure 23-5, a group is displayed with a two-headed icon (for example, the Domain Admins group). Single users are displayed with a single head, as is the case with the Administrator. By grouping users together, you can save a lot of

time when assigning permissions to files and other resources; instead of assigning permissions to one user at a time, you can assign permissions to the entire group in one shot.

Account permissions such as file and printer access can be assigned to individual users or to groups. These permissions (also known as access modes) are examples of access control lists (ACLs)—specifically, file system access control lists (abbreviated as FACL or FSACL). An ACL is a list of permissions attached to an object. ACLs reside on firewalls, routers, and computers. Permissions in an ACL might allow access or deny access. The configuration depends on who is required to have access; in that case, the configuration is up to you.

NOTE In a Windows environment, ACLs are broken down into individual access control entries (ACEs); for example, the user JohnT has Read access to the marketing folder. Also, Microsoft uses the term *discretionary access control list (DACL)* to refer to an ACL that identifies trustees who are allowed or denied access to an object. It uses the term *system access control list (SACL)* to refer to an ACL that enables administrators to log attempts to access a secured object. These are examples of the ***access policies*** utilized in a Windows environment.

In Windows, there are two types of permissions. Sharing permissions are basic permissions including Full Control, Change, and Read, which are applied to folders only. These are often ignored in favor of the more powerful (and superseding) NTFS permissions, also called security permissions, which can secure folders and individual files. In a standard Windows folder on a domain, the types of NTFS permissions include the following:

- Full Control
- Modify
- Read & Execute
- List Folder Contents
- Read
- Write

These permissions are shown in Figure 23-8 on a Windows Server in the Properties window of a folder named testfolder. Note that the Administrators group has full control of the folder. Also note that you can allow particular permissions or specifically deny those permissions. If a permission is not set to Allow, it is implicitly denied.

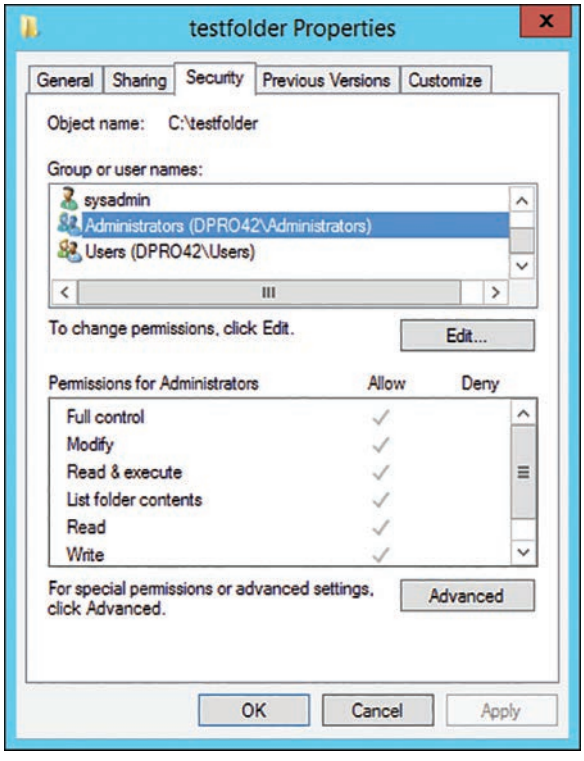


FIGURE 23-8 NTFS Permissions

In Linux, file permissions are broken down into three types: read, write, and execute (R, W, and X). They can be assigned to three different permission groups: owner, group, and all users (U, G, and O or A). They can be assigned and configured at the command line with the **chmod** command (change mode), either by group letter or by using a designated numbering system. For example, in the latter case:

- R (Read) = 4
- W (Write) = 2
- X (Execute) = 1

These are added together to form each permission. For example, if a user was given read and write access, the number would be 6. If, however, the person was given read, write, and execute access, the number would be 7. Here's an example of the **chmod** command:

```
chmod 760 testfile
```

Here, the **chmod** command is followed by the numbers 7, 6, and 0, and the name of the file, testfile (of course, the path to that file can be more complex). The first number, 7, represents the owner permission (which in this case is equal to full access). The second number, 6, represents the group permission (in this case, read and write access). The final number, 0, represents the all users (or all *other* users) permission, which has no access. This is just an example; the numbers could be whatever you select for the three groups. You might see 777 when all groups have all permissions, though it is not normally recommended. It is common on a web server or file server to see 755, which means that the owner has all permissions, and the group and all users have read and execute permissions. You don't want a typical user to have write permissions on a web server. In summary, the **chmod** command is a great way to assign permissions to all three groups using a single command.

To avoid access violations when working with permissions, you should implement the least privilege or minimal privilege concept. You should give the users only the amount of access that they absolutely need. Note that permissions for long-term employees could suffer from *privilege creep* over time. To mitigate this, you should consider periodic user permission reviews and evaluation of ACLs. This permission auditing procedure ensures that users have access to the correct data. In general, this is known as *user access recertification*. Consider this procedure if a company has a particularly high attrition rate (hiring and terminating of employees). This verifies that users no longer with the company cannot log on to the network and cannot gain access to resources. It also ensures that new users can gain access to necessary resources.

Permission Inheritance and Propagation

If you create a folder, the default action it takes is to inherit permissions from the parent folder, which ultimately come from the root folder. So any permissions set in the parent are inherited by the subfolder. To view an example, locate any folder within an NTFS volume (besides the root folder), right-click it, and select Properties, access the Security tab, and click the Advanced button. That brings up a window, as shown in Figure 23-9. At the bottom of the figure, you can see whether the folder is inheriting permissions from the parent. In this case, it is inheriting permissions.

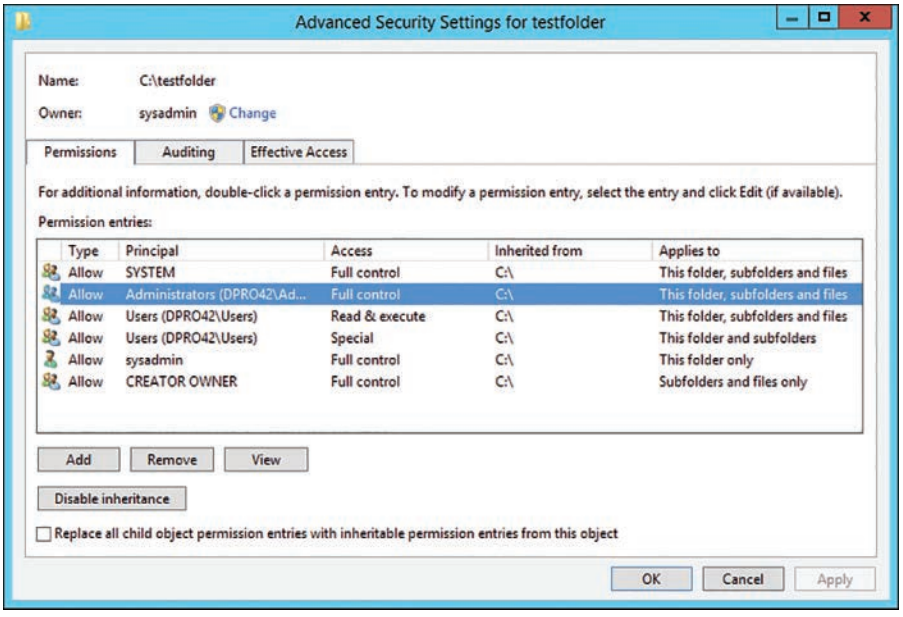


FIGURE 23-9 Inheritable Permissions

Here, any permissions added or removed in the parent folder are also added or removed in the current folder. In addition, those permissions inherited cannot be modified in the current folder. To make modifications, you would have to disable inheritance, either by clicking the button shown in Figure 23-9 or by deselecting a checkbox on older versions of Windows Server. When you do so, you have the option to copy the permissions from the parent to the current folder or remove them entirely. To summarize, by default, the parent automatically propagates permissions to the subfolder, and the subfolder inherits its permissions from the parent.

You can also propagate permission changes to subfolders not inheriting from the current folder. To do so, select the Replace All Child Object Permission Entries checkbox.

This process might seem a bit confusing, and you will probably not be asked many questions on the subject. Just remember that folders automatically inherit from the parent unless you turn off inheriting, and you can propagate permission entries to subfolders at any time by selecting the Replace option.

NOTE In a Windows environment, there are two types of permissions: share permissions and NTFS permissions. By default, the more restrictive of the two sets of permissions is applied to the user. However, quite often you can configure NTFS permissions to take precedence over share permissions—and effectively "ignore" the share permissions.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 23-5 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 23-5 Key Topics for Chapter 23

Key Topic Element	Description	Page Number
Section	Identity	623
Table 23-2	Authentication Methods	626
Section	Account Policies	633
Figure 23-4	Identity and Access Lifecycle	634
Section	Password Management	636
List	Requirements when building policies, processes, and standards around password creation	636
Table 23-3	Password-Generation Methods	637
Table 23-4	Attributes and Identity Terms Used in Account Policies	639

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

identity, identity provider (IdP), token, smart card, certificate, SSH keys, biometric attribute, guest accounts, shared and generic accounts, user account, service account, least privilege, attribute-based access control (ABAC), geofencing, account permissions, access policies

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. Which command can be used to determine who is logged in to a Linux system?
2. What factor of multifactor authentication utilizes something you are?
3. How many factors are needed for multifactor authentication?
4. Which type of password is generated by an external entity and synchronized with internal resources?
5. What type of access control is dynamic and context-aware?
6. _____ is the actual method of determining the physical location of the user trying to authenticate.

This page intentionally left blank



This chapter covers the following topics related to Objective 3.8 (Given a scenario, implement authentication and authorization solutions) of the CompTIA Security+ SY0-601 certification exam:

- Authentication management
  - Password keys
  - Password vaults
  - TPM
  - HSM
  - Knowledge-based authentication
- Authentication/authorization
  - EAP
  - Challenge-Handshake Authentication Protocol (CHAP)
  - Password Authentication Protocol (PAP)
  - 802.1X
  - RADIUS
  - Single sign-on (SSO)
  - Security Assertions Markup Language (SAML)
  - Terminal Access Controller Access Control System Plus (TACACS+)
  - OAuth
  - OpenID
  - Kerberos
- Access control schemes
  - Attribute-based access control (ABAC)
  - Role-based access control
  - Rule-based access control
  - MAC
  - Discretionary access control (DAC)
  - Conditional access
  - Privileged access management
  - Filesystem permissions

Implementing Authentication and Authorization Solutions

Authentication and authorization solutions are among the fundamental building blocks of an IT organization, and obviously they are a key control when it comes to security. This chapter starts with an overview of authentication management, including a discussion of password keys, password vaults, TPM, HSM, and knowledge-based authentication concepts. From there we dig deeper into the specifics of authentication and authorization with topics such as Extensible Authentication Protocol (EAP) and Challenge-Handshake authentication protocol (CHAP), including concepts such as Password Authentication Protocol (PAP), 802.1X, RADIUS, single sign-on (SSO), Security Assertions Markup Language (SAML), and Terminal Access Controller Access Control System Plus (TACACS+). Other topics that are touched on include OAuth, OpenID and Kerberos. We finish the chapter learning about access control schemes.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 24-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 24-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Authentication Management	1–3
Authentication/Authorization	4–6
Access Control Schemes	7–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is a type of technology used by inserting a hardware USB device into a computer for authentication?
   a. Password key
   b. Password vault
   c. Trusted Platform Module
   d. None of these answers are correct.

2. Which of the following can be used to store a set of credentials for later use?
   a. Password key
   b. Password vault
   c. Password authentication module
   d. All of these answers are correct.

3. Which of the following is utilized to verify a user based on something only that person knows?
   a. Knowledge-based authentication
   b. Password vault
   c. Password key
   d. Hardware security module

4. Which of the following is a system that is utilized to allow a user to log in once but gain access to multiple systems without being asked to log in again?
   a. Password key
   b. Password vault
   c. Two-factor authentication
   d. Single sign-on

5. Which of the following is a system where a user's identity and attributes are shared across multiple identity management systems?
   a. Trusted Platform Module
   b. Knowledge-based authentication
   c. Password authentication
   d. Federated identity management

6. Which of the following is an open standard for exchanging authentication and authorization data between identity providers?
   a. KBA
   b. SAML
   c. SSO
   d. None of these answers are correct.

7. Which of the following is an authentication scheme where the access control policy is generally determined by the owner?
   a. MAC
   b. DAC
   c. RBAC
   d. SAML

8. Which of the following is an authentication scheme where the access control policy is determined by a computer system, not by a user or owner?
   a. 802.1X
   b. MAC
   c. RBAC
   d. SAML

9. Which of the following is a security principle where users are given only the privileges needed to do their job?
   a. Single sign-on
   b. Least privilege
   c. Knowledge-based authentication
   d. All of these answers are correct.

10. Which of the following is an access model that works with sets of permissions instead of individual permissions that are label-based?
    a. MAC
    b. RBAC
    c. DAC
    d. ABAC

Foundation Topics

Authentication Management

Key Topic

Authentication is one of the primary elements to address when securing your computing environment. It can also be one of the most challenging aspects for end users. For this reason, a number of technologies have emerged to address these challenges. In the following sections, we discuss some of them, including password keys, password vaults, Trusted Platform Modules, and hardware security modules.

Password Keys

Password keys are a technology typically deployed by corporations when implementing two-factor authentication. The primary use case is remote access to the organization's environment. However, many organizations also use them internally. These keys are especially important to use for access to highly sensitive data or applications that serve that data—for instance, financial applications or anything involving intellectual property, such as source code.

Many types of password keys are on the market these days. They come in various form factors and are utilized in different ways. For instance, some are used by inserting into a computer USB port. Some are simply one-time password tokens used as a second factor when authenticating, whereas others are a combination of different functions such as those just mentioned.

Password Vaults

One of the challenges with traditional user/password authentication is remembering the different sets of credentials. Most people have more than 25 different sets of credentials that they need to remember. If these users are following best practices, they should have a different username and password for each account. Every new email or website that a user creates an account with adds another set of credentials to the list. A ***password vault*** helps solve the issue of credential storage by providing a central system that stores the various sets of credentials in a secure management system. The vault then has its own set of credentials and possibly another authentication factor that is used to access the vault when the credentials are needed for accessing a system. A password vault is also often referred to as a *password manager*. It is simply a piece of software that is utilized to store and manage credentials. Typically, the credentials are stored in an encrypted database. Having an encrypted database protects the credentials from being compromised if the database is obtained by a threat actor through the compromise of a system holding the database.

Trusted Platform Module

The ***Trusted Platform Module (TPM)*** is a chip residing on the motherboard that stores the encrypted keys. NIST defines a TPM as a tamper-resistant integrated circuit built into some computer motherboards that can perform cryptographic operations (including key generation) and protect small amounts of sensitive information, such as passwords and cryptographic keys. The security mechanisms that make the TPM tamper resistant are utilized to prevent a threat actor or malware from being able to tamper with the functions of the Trusted Platform Module. A TPM is a form of hardware security module.

Hardware Security Modules

Hardware security modules (HSMs) are physical devices that act as secure cryptoprocessors. This means that they are used for encryption during secure login/authentication processes, during digital signings of data, and for payment security systems. The beauty of a hardware-based encryption device such as an HSM (or a TPM) is that it is faster than software encryption.

HSMs can be found in adapter card form, as devices that plug into a computer via USB, and as network-attached devices. They are generally tamper-proof, giving a high level of physical security. They can also be used in high-availability clustered environments because they work independently of other computer systems and are used solely to calculate the data required for encryption keys. However, many of these devices require some kind of management software to be installed on the computer they are connected to. Some manufacturers offer this software as part of the purchase, but others do not, forcing purchasers to build the management software themselves. Due to this lack of management software and the cost involved in general, HSMs have seen slower deployment with some organizations. This concept also holds true for hardware-based drive encryption solutions.

Often, HSMs are involved in the generation, storage, and archiving of encrypted key pairs such as the ones used in Transport Layer Security (TLS) sessions online, public key cryptography, and public key infrastructures (PKIs).

Knowledge-Based Authentication

NIST defines ***knowledge-based authentication (KBA)*** as authentication of an individual based on knowledge of information associated with his or her claimed identity in public databases. Knowledge of such information is considered to be private rather than secret, because it may be used in contexts other than authentication to a verifier, thereby reducing the overall assurance associated with the authentication process.

A popular use case for this type of authentication is to recover a username or reset a password. Typically, a set of predetermined questions is asked of the user. These questions must have already been provided by the end user at the time of account setup or provided as an authenticated user at a later time. The idea is that the information that was provided is something only the user would know. That is why it is important to utilize a set of questions that cannot be easily guessed or are not public knowledge.

Key Topic Authentication/Authorization

Now that we've covered some physical authentication methods, let's move into authentication models, components, and technologies used to grant or deny access to operating systems and computer networks.

The first thing you, as security administrator, should do is plan what type of authentication model to use. Then you should consider what type of authentication technology and how many factors of authentication will be implemented. Also, you should consider how the authentication system will be monitored and logged. Getting more into the specifics, will only local authentication be necessary? Or will remote authentication also be needed? And which type of technology should be utilized? Will it be Windows-based or a third-party solution? Let's look at these concepts now and examine some different examples of the possible solutions you can implement.

Many small businesses and even some midsized businesses often have one type of authentication to gain access to a computer network: the username and password. In today's security-conscious world, having only this one type of authentication is not enough for the average organization. Some companies share passwords or fail to enforce password complexity. In addition, password-cracking programs are becoming more and more powerful and work much more quickly than they did just a few years ago, making the username and password authentication scheme limiting. We're not saying that this approach shouldn't be used, but perhaps it should be enforced, enhanced, and integrated with other technologies.

Because of the limitations of a single type of authentication such as username and password, organizations sometimes use multiple factors of authentication. In *multifactor authentication (MFA)*, two or more types of authentication are used for user access control. One example of multifactor authentication would be when a user needs to sign in with a username and password and swipe some type of smartcard or use some other type of physical token at the same time. Adding factors of authentication makes it more difficult for a malicious person to gain access to a computer network or an individual computer system. Sometimes an organization uses three factors of authentication—perhaps a smartcard, biometrics, and a

username/password. The disadvantages of a multifactor authentication scheme are that users need to remember more information and remember to bring more identification with them, and more IT costs and more administration are involved. Another disadvantage of some MFA environments is that they are static: rules and whitelists/blacklists are usually configured manually.

A more dynamic way of authenticating individuals is to utilize *context-aware authentication* (also known as context-sensitive access). It is an adaptive way of authenticating users based on their usage of resources and the confidence that the system has in these users. This approach can automatically increase the level of identification required and/or increase or decrease the level of access to resources based on constant analysis of the users.

In some organizations an individual user might need access to several computer systems. By default, each of these systems has a separate login. It can be difficult for users to remember the various logins. Using ***single sign-on (SSO)***, a user can log in once but gain access to multiple systems without being asked to log in again. This system is complemented by single sign-off, which is basically the reverse; logging off signs off a person from multiple systems. Single sign-on is meant to reduce password fatigue or password chaos, which occurs when a person can become confused and possibly even disoriented when having to log in with several different usernames and passwords. It is also meant to reduce IT help desk calls and password resets. By implementing a more centralized authentication scheme such as single sign-on, many companies have reduced IT costs significantly. If implemented properly, single sign-on can also reduce phishing. In large networks and enterprise scenarios, it might not be possible for users to have a single sign-on, and in these cases, the security principle might be referred to as *reduced* sign-on. Single sign-on can be Kerberos-based, integrated with Windows authentication, or token- or smartcard-based.

SSO is a derivative of *federated identity management* (also called FIM or FIdM). In this system, a user's identity and attributes are shared across multiple identity management systems. These various systems can be owned by one organization; for example, Microsoft offers the Forefront Identity Manager software, which can control user accounts across local and cloud environments. Also, Google, Yahoo!, and Amazon utilize this federation approach. But some providers join forces so that information can be shared across multiple services and environments between the companies, yet still allow the user a single login. Shibboleth is an example of an SSO system that allows people to sign in with a single digital identity and connect to various systems run by federations of different organizations. SSO systems—and federated systems in general—often incorporate the concept of transitive trust where two networks (or more) have a relationship such that users logging in to one network get access to data on the other.

While an SSO is easier for users to remember, it also acts as a single point of failure. In addition, sometimes a company might not be watching out for the users' best interests—either unwittingly or otherwise—and might fail to realize that multiple systems have been configured as a transitive trust. Let's say that a user has an account with Company A and has a separate account with Company B. Imagine that Companies A and B have a two-way trust. Now, let's say there is a third organization, Company C, that has a two-way trust with Company B. At this point, the user's account information from Companies A and B could be shared with Company C even though the user never signed up with that company. This kind of activity is frowned upon, but the user might not even know when it happens: two companies might merge, or a company might be bought out or otherwise absorbed by another. So, when it comes to authentication, it is sometimes wise to avoid trust relationships and strongly consider whether single sign-on will ultimately be more beneficial or costly to your organization.

Another concern is web-based SSO, which can be problematic due to disparate proprietary technologies. To help alleviate this problem, the XML-based Security Assertion Markup Language (SAML) and the OpenID Connect protocol were developed. OpenID Connect is an interoperable authentication protocol based on the OAuth 2.0 family of specifications. It uses straightforward REST/JSON message flows with a design goal of "making simple things simple and complicated things possible." Both OpenID Connect and SAML specify separate roles for the user, the service provider, and the identity provider. Shibboleth is also based on SAML.

Security Assertion Markup Language

Security Assertion Markup Language (SAML) is an open standard for exchanging authentication and authorization data between identity providers. SAML is used in many SSO implementations.

The OASIS Security Assertion Markup Language standard is currently the most used standard for implementing federated identity processes. SAML is an XML-based framework that describes the use and exchange of SAML assertions in a secure way between business entities. The standard describes the syntax and rules to request, create, use, and exchange these assertions.

The SAML process involves a minimum of two entities: the *SAML assertion party* (or *SAML authority*), which is the entity that produces the assertion, and the *SAML relying party*, which is the entity that uses the assertion to make access decisions.

An assertion is the communication of security information about a subject (also called a principal) in the form of a statement. The basic building blocks of SAML are the SAML assertion, SAML protocol, SAML binding, and SAML profile. SAML assertions can contain the following information:

- **Authentication statement:** Includes the result of the authentication and additional information such as the authentication method, timestamps, and so on
- **Attribute statement:** Includes attributes about the principal
- **Authorization statement:** Includes information on what the principal is allowed to do

An example of an assertion would be User A, who has the email address usera@domain.com authenticated via username and password, is a platinum member, and is authorized for a 10 percent discount.

SAML protocols define the protocols used to transfer assertion messages. SAML bindings include information on how lower-level protocols (such as HTTP or SOAP) transport SAML protocol messages. SAML profiles are specific combinations of assertions, protocols, and bindings for specific use cases. Examples of profiles include Web Browser Single Sign-On, Identity Provider Discovery, and Enhanced Client and Proxy (ECP).

Figure 24-1 shows the SAML building blocks.

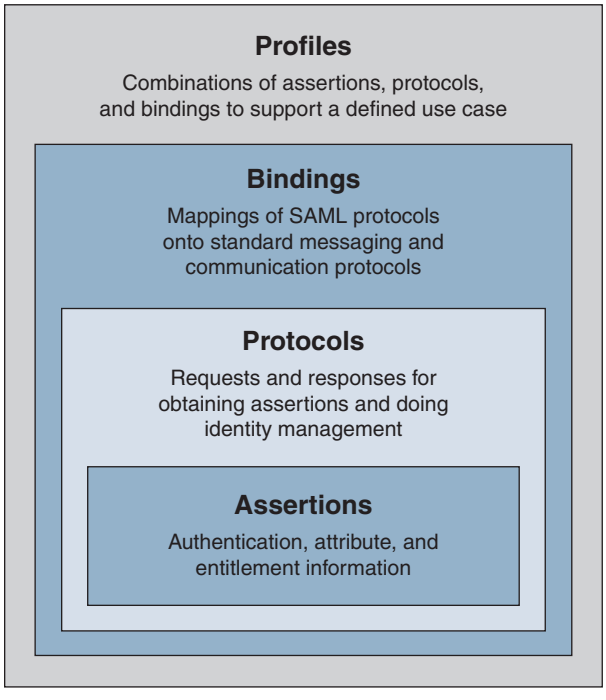


FIGURE 24-1 SAML Building Blocks

SAML also defines the concepts of identity provider (IdP) and service provider (SP). SAML can work in two different ways:

- In IdP-initiated mode, a user is already authenticated on the IdP and requests a service from the SP (for example, by clicking a link on the IdP website). The IdP builds an assertion that is sent to the SP within the user request to the SP itself.

  For example, a user who is authenticated on an airline website decides to book a rental car by clicking a link on the airline website. The airline IAM system, which assumes the role of an IdP, sends assertion information about the user to the rental car IAM, which in turn authenticates the user and provides access rights based on the information in the assertion.

- In SP-initiated mode, a user initiates an access request to some resource on the SP. Because the federated identity is managed by a different IdP, the SP redirects the user to log in at the IdP. After the login, the IdP sends a SAML assertion back to the SP.

Figure 24-2 shows an example of IdP-initiated mode (on the left) and SP-initiated mode (on the right).

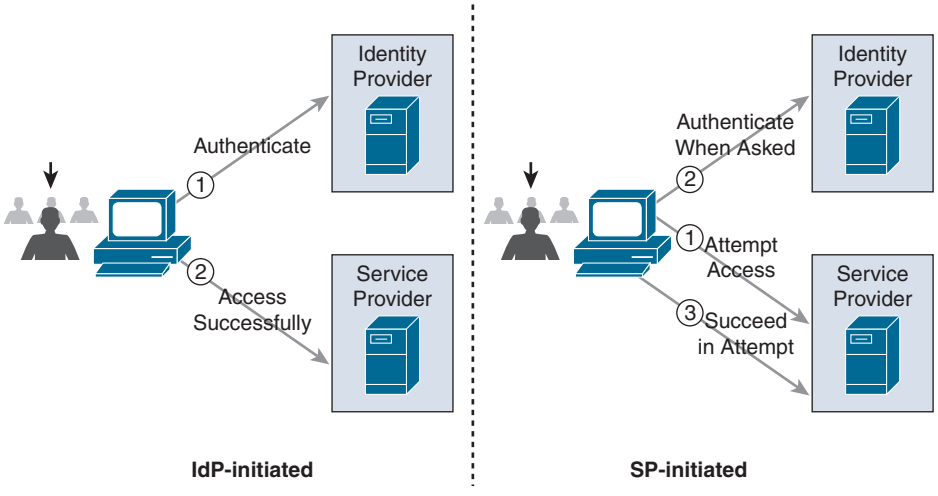


FIGURE 24-2 SAML IdP-Initiated Mode and SP-Initiated Mode

OAuth

OAuth is a framework that provides authorization to a third-party entity (for example, a smartphone application) to access resources hosted on a resource server. In a classic client/server authorization framework, the third-party entity would receive

the credentials from the resource owner (user) and then access the resource on the resource server.

The main issue OAuth resolves is providing the third-party entity authorization to access restricted resources without passing the client credentials to this third party. Instead of getting the user credentials, the entity requesting access receives an authorization token that includes authorization information, such as scope and duration, and that is used to request access to a resource hosted by the resource server. The OAuth schema is usually called *delegation of access*.

OAuth 2.0, defined in RFC 6749, includes four main roles:

- **Resource owner:** The party that owns the resource (for example, a user) and grants authorization to access some of its resources
- **Client:** The party that requires access to a specific resource
- **Resource server:** The party that hosts or stores the resource
- **Authorization server:** The party that provides an authorization token

In the basic scenario, the authorization is done with six messages:

1. The client sends an authorization request to the resource owner or indirectly to the authorization server.
2. The resource owner (or the authorization server on behalf of the resource owner) sends an authorization grant to the client.
3. The client sends the authorization grant to the authorization server as proof that authorization was granted.
4. The authorization server authenticates the client and sends an access token.
5. The client sends the access token to the resource server as proof of authentication and authorization to access the resources.
6. The resource server validates the access token and grants access.

For example, a user (the resource owner) may grant access to personal photos hosted at some online storage provider (the resource server) to an application on a mobile phone (the client) without directly providing credentials to the application but instead by directly authenticating with the authorization server (in this case, also the online storage provider) and authorizing the access.

Figure 24-3 shows an example of an OAuth exchange.

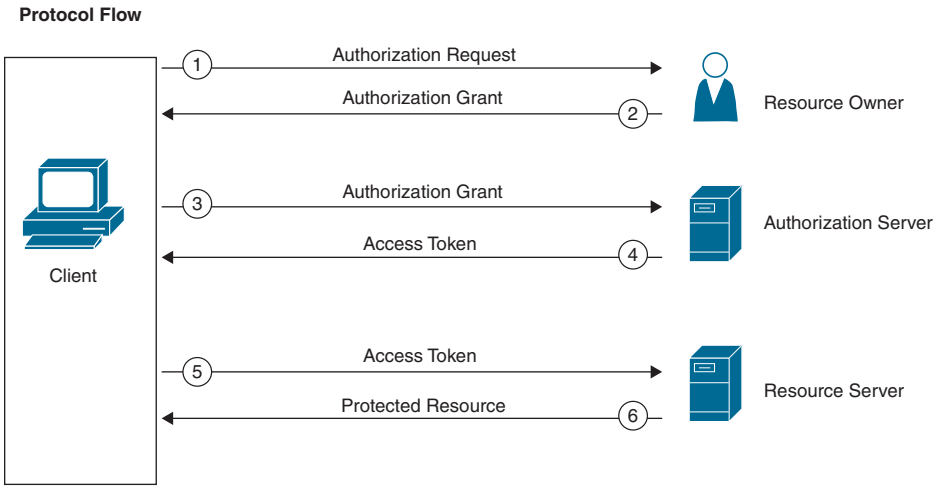


FIGURE 24-3 OAuth Exchange

OpenID and OpenID Connect

OpenID has been a popular SSO protocol for federated systems for quite some time. In the 2.0 version, the authentication and authorization process is similar to the one in SAML. OpenID also defines an IdP, called the OpenID provider (OP), and a relying party (RP), which is the entity that holds the resource the user wants to access. In OpenID, a user is free to select an OP of his or her choice, and the initial identity is provided in the form of a URL.

Version 2.0 has been superseded by OpenID Connect. This version drops the authorization functionality present in version 2.0 and is designed to work with OAuth 2.0 for deployments. In practice, OpenID Connect operates as an authentication profile for OAuth. In OpenID Connect, when a user tries to access resources on an RP, the RP sends an authentication request to the OP for that user. In practice, this is an OAuth 2.0 authorization request to access the user's identity at the OP. The authentication request can be of three types:

- Authorization code flow (the most commonly used)
- Implicit flow
- Hybrid flow

In an authorization code flow scenario, after the user authenticates with the OP, the OP asks the user for consent and issues an authorization code that the user then sends to the RP. The RP uses this code to request an ID token and access token from the OP, which is the way the OP provides assertion to the RP.

Whatever the type of authentication scheme used, it needs to be monitored periodically to make sure that it's working properly. The authentication system should block people who cannot furnish proper identification and should allow access to people who do have proper identification.

NOTE Keep in mind that OpenID is not the same as OpenID Connect. OpenID is an open standard and decentralized authentication protocol controlled by the OpenID Foundation. OAuth is an open standard for access delegation. OpenID Connect (OIDC) combines the features of OpenID and OAuth and provides both authentication and authorization.

Key Topic

802.1X and EAP

802.1X is an IEEE standard that defines port-based network access control (PNAC). Not to be confused with 802.11x WLAN standards, 802.1X is a data link layer authentication technology used to connect hosts to a LAN or WLAN. 802.1X enables you to apply a security control that ties physical ports to end-device MAC addresses and prevents additional devices from being connected to the network. It is a good way of implementing port security, much better than simply setting up MAC filtering.

It all starts with the central connecting device such as a switch or wireless access point. These devices must first enable 802.1X connections; they must have the 802.1X protocol (and supporting protocols) installed. Vendors that offer 802.1X-compliant devices (for example, switches and wireless access points) include Cisco, Symbol Technologies, and Intel. Next, the client computer needs to have an operating system, or additional software, that supports 802.1X. The client computer is known as the supplicant. All recent Windows versions support 802.1X. macOS offers support as well, and Linux computers can use Open1X to enable client access to networks that require 802.1X authentication.

802.1X encapsulates the ***Extensible Authentication Protocol (EAP)*** over wired or wireless connections. EAP is not an authentication mechanism in itself but instead defines message formats. 802.1X is the authentication mechanism and defines how EAP is encapsulated within messages. Figure 24-4 shows an example of an 802.1X-enabled network adapter. In the figure, you can see that the box for enabling 802.1X has been checked and that the type of network authentication method for 802.1X is EAP—specifically, Protected EAP (PEAP).

Key Topic

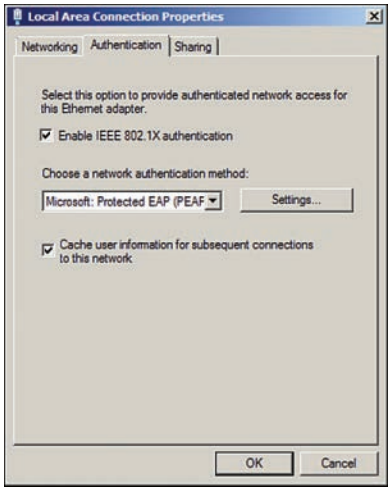


FIGURE 24-4 An 802.1X-Enabled Network Adapter in Windows

NOTE You can enable 802.1X in Windows by accessing the Local Area Connection Properties page.

Following are three components to an 802.1X connection:

- **Supplicant:** A software client running on a workstation. This is also known as an authentication agent.
- **Authenticator:** A wireless access point or switch.
- **Authentication server:** An authentication database, most likely a RADIUS server.

The typical 802.1X authentication procedure has four steps:

Step 1. **Initialization:** If a switch or wireless access point detects a new supplicant, the port connection enables port 802.1X traffic; other types of traffic are dropped.

Step 2. **Initiation:** The authenticator (switch or wireless access point) periodically sends EAP requests to a MAC address on the network. The supplicant listens for this address and sends an EAP response that might include a user ID or other similar information. The authenticator encapsulates this response and sends it to the authentication server.

Step 3. **Negotiation:** The authentication server then sends a reply to the authenticator. The authentication server specifies which EAP method to use. (These are listed next.) Then the authenticator transmits that request to the supplicant.

Step 4. **Authentication:** If the supplicant and authentication server agree on an EAP method, the two transmit until there is either success or failure to authenticate the supplicant computer.

Figure 24-5 illustrates the components used in these steps.

Key Topic

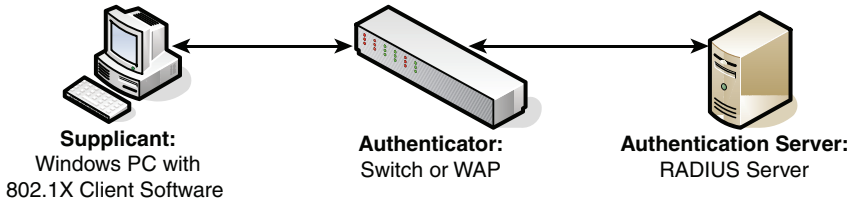


FIGURE 24-5 Components of a Typical 802.1X Authentication Procedure

Following are several types of EAP authentication:

- **EAP-MD5:** This challenge-based authentication provides basic EAP support. It enables only one-way authentication and not mutual authentication.
- **EAP-TLS:** This version uses Transport Layer Security, which is a certificate-based system that does enable mutual authentication. It does not work well in enterprise scenarios because certificates must be configured or managed on the client side and server side.
- **EAP-TTLS:** This version is Tunneled Transport Layer Security and is basically the same as TLS except that it is done through an encrypted channel, and it requires only server-side certificates.
- **EAP-FAST:** This version uses a protected access credential instead of a certificate to achieve mutual authentication. FAST stands for Flexible Authentication via Secure Tunneling.
- **PEAP:** This is the Protected Extensible Authentication Protocol (also known as Protected EAP). It uses MS-CHAPv2, which supports authentication via Microsoft Active Directory databases. It competes with EAP-TTLS and includes legacy password-based protocols. It creates a TLS tunnel by acquiring a public key infrastructure (PKI) certificate from a server known as a certificate authority (CA). The TLS tunnel protects user authentication much like EAP-TTLS.

Cisco also created a proprietary protocol called Lightweight EAP (LEAP), and it is just that—proprietary. To use LEAP, you must have a Cisco device such as an Aironet WAP or Catalyst switch, or another vendor's device that complies with the Cisco Compatible Extensions program. Then you must download a third-party client on Windows computers to connect to the Cisco device. Most WLAN vendors offer an 802.1X LEAP download for their wireless network adapters.

Although 802.1X is often used for port-based network access control on the LAN, especially VLANs, it can also be used with virtual private networks (VPNs) as a way of remote authentication. Central connecting devices such as switches and wireless access points remain the same, but on the client side 802.1X would need to be configured on a VPN adapter instead of a network adapter.

Many vendors, such as Intel and Cisco, refer to 802.1X with a lowercase *x*; however, the IEEE displays this on its website with an uppercase *X*, as does the IETF.

LDAP

The Lightweight Directory Access Protocol (LDAP) is an application layer protocol used for accessing and modifying directory services data. It is part of the TCP/IP suite. Originally used in WAN connections, it has developed over time into a protocol commonly used by services such as Microsoft Active Directory on Windows server domain controllers. LDAP acts as the protocol that controls the directory service. This service organizes the users, computers, and other objects within the Active Directory. Figure 24-6 shows an example of the Active Directory. Take note of the highlighted list of users (known as objects of the Active Directory) from the Users folder. Also, observe other folders such as Computers that house other objects (such as Windows client computers).

Key Topic

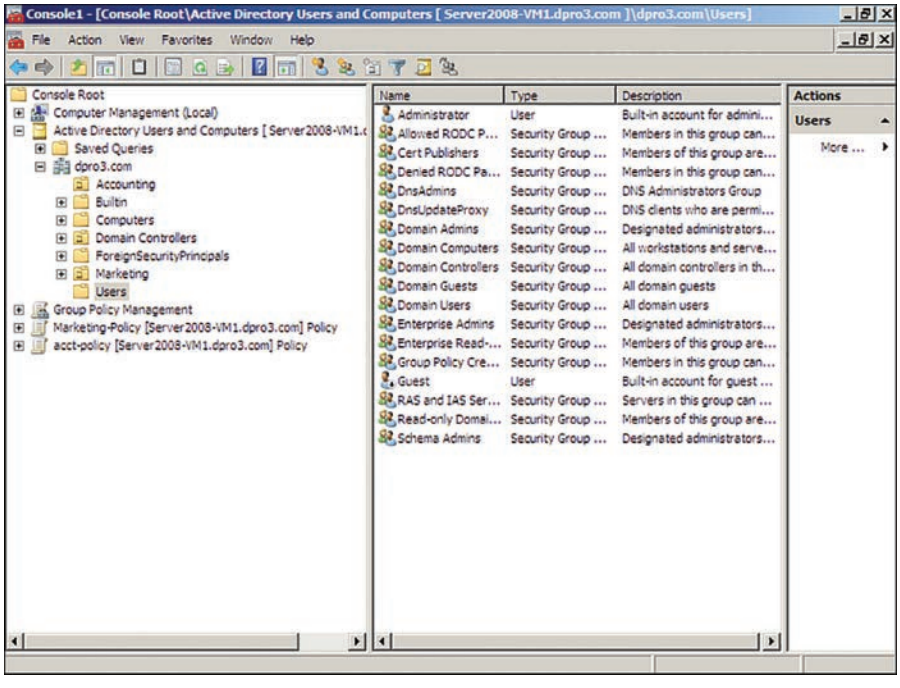


FIGURE 24-6 Active Directory Showing User Objects

NOTE Windows servers running Active Directory use parameters and variables when querying the names of objects—for example, CN=dprowse, where *CN* stands for common name and *dprowse* is the username. Taking it to the next level, consider DC=ServerName. *DC* stands for domain component, and *ServerName* is the variable and the name of the server. Microsoft is famous for using the name *fabrikam* as its test name, but, of course, you would use whatever your server name is. In the case of fabrikam, an entire LDAP query might look something like this:

```
<LDAP://DC=Fabrikam, DC=COM>
```

A Microsoft server that has Active Directory and LDAP running has inbound port 389 open by default. To protect Active Directory from being tampered with, you can use Secure LDAP, also known as LDAPS. It was covered previously in Chapter 17, "Implementing Secure Protocols."

Kerberos and Mutual Authentication

Kerberos is an authentication protocol designed at MIT; it enables computers to prove their identity to each other in a secure manner. It is used most often in a client/server environment; the client and server both verify each other's identity. This is known as two-way authentication or mutual authentication. Often, Kerberos protects a network server from illegitimate login attempts, just as the mythological three-headed guard dog of the same name (also known as Cerberus) guards Hades.

A common implementation of Kerberos occurs when a user logs on to a Microsoft domain. (Of course, we are not saying that Microsoft domains are analogous to Hades!) The domain controller in the Microsoft domain is known as the key distribution center (KDC). This server works with *tickets* that prove the identity of users. The KDC is composed of two logical parts: the authentication server and the ticket-granting server. Basically, a client computer attempts to authenticate itself to the authentication server portion of the KDC. When it does so successfully, the client receives a ticket. This is actually a ticket to get other tickets—known as a ticket-granting ticket (TGT). The client uses this preliminary ticket to demonstrate its identity to a ticket-granting server in the hopes of ultimately getting access to a service—for example, making a connection to the Active Directory of a domain controller.

The domain controller running Kerberos has inbound port 88 open to service logon requests from clients. Figure 24-7 shows a **netstat -an** command run on a Windows server that has been promoted to a domain controller. It points out port 88 (used by Kerberos) and port 389 (used by LDAP) on the same domain controller.

Key Topic

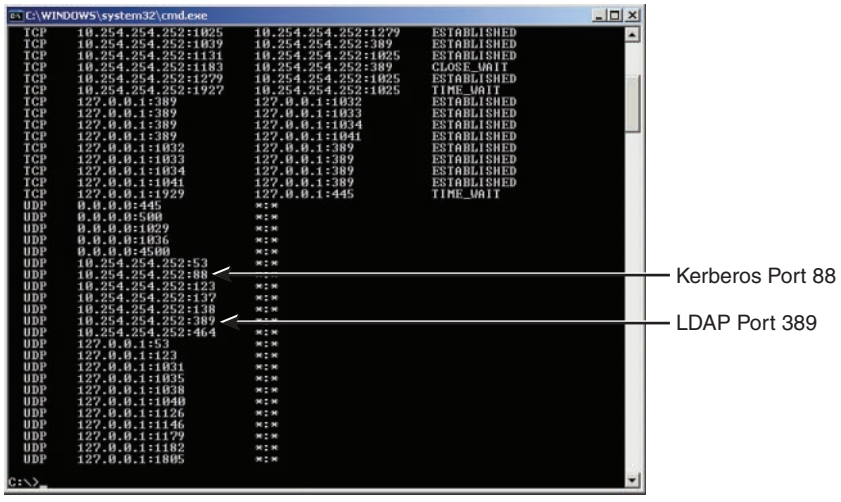


FIGURE 24-7 Results of the **netstat -an** Command on a Windows Server

Kerberos is designed to protect against replay attacks and eavesdropping. One of the drawbacks of Kerberos is that it relies on a centralized server such as a domain controller. This can be a single point of failure. To alleviate this problem, you can install secondary and tertiary domain controllers that keep a copy of the Active Directory and are available with no downtime in case the first domain controller fails.

Another possible issue is one of synchronicity. Time between the clients and domain controller must be synchronized for Kerberos to work properly. If, for some reason, a client attempting to connect to a domain controller becomes desynchronized, it cannot complete the Kerberos authentication, and as a result, the user cannot log on to the domain. This situation can be fixed by logging on to the affected client locally and synchronizing the client's time to the domain controller by using the **net time** command. For example, to synchronize to the domain controller in Figure 24-7, the command would be

```
net time \\10.254.254.252 /set
```

Afterward, the client should be able to connect to the domain.

Kerberos—like any authentication system—is vulnerable to attack. Older Windows operating systems that run, or connect to, Kerberos are vulnerable to privilege escalation attacks; and newer Windows operating systems are vulnerable to spoofing. Of course, Microsoft quickly releases updates for these kinds of vulnerabilities (as they are found), but if you do not allow Windows Update to automatically update, it's important to review the CVEs for the Microsoft systems often.

NOTE In a Red Hat Enterprise environment that uses an SSO such as Kerberos, *pluggable authentication modules (PAMs)* can be instrumental in providing both the system administrator and security administrator with flexibility and control over authentication, as well as fully documented libraries for the developer.

Remote Authentication Technologies

Even more important than authenticating local users is authenticating *remote* users. The chances of illegitimate connections increase when you allow remote users to connect to your network. Examples of remote authentication technologies include RAS, VPN, RADIUS, TACACS+, and CHAP.

Remote Access Service

Remote Access Service (RAS) began as a service that enabled dial-up connections from remote clients. Nowadays, more and more remote connections are made with high-speed Internet technologies such as cable Internet and fiber-optic connections. But you can't discount the dial-up connection. It is used in certain areas where other Internet connections are not available and is still used as a fail-safe in many network operation centers and server rooms to take control of networking equipment.

One of the best things you can do to secure an RAS server is to deny access to individuals who don't require it. Even if the user or user group is set to "not configured," it is wise to specifically deny them access. You should allow access to only those users who need it and on a daily basis monitor the logs that list who connected. If there are any unknowns, you should investigate immediately. Also, be sure to update the permissions list often in the case that a remote user is terminated or otherwise leaves the organization.

The next most important security precaution is to set up RAS authentication. One secure way is to use the ***Challenge-Handshake Authentication Protocol (CHAP)***, which is an authentication scheme used by the Point-to-Point Protocol (PPP), which in turn is the standard for dial-up connections. It uses a challenge-response mechanism with one-way encryption. Therefore, it is not capable of mutual authentication in the way that Kerberos is. Microsoft developed its own version of CHAP, known as MS-CHAP; an example is shown in Figure 24-8. The figure shows the Advanced Security Settings dialog box of a dial-up connection. In this particular configuration, notice that encryption is required and that the only protocol allowed is MS-CHAPv2. It's important to use version 2 of MS-CHAP because it provides for mutual authentication between the client and authenticator. Of course, the RAS server has to be configured to accept MS-CHAP connections as well. You also have the option to enable EAP for the dial-up connection. Other RAS authentication protocols include SPAP, which is of lesser security, and ***Password Authentication***

Protocol (PAP), which sends usernames and passwords in clear text—obviously insecure and to be avoided. Therefore, PAP should not be used in networks today.

NOTE MS-CHAP utilizes the MD5 hashing algorithm, which is known to be insecure.

Key Topic

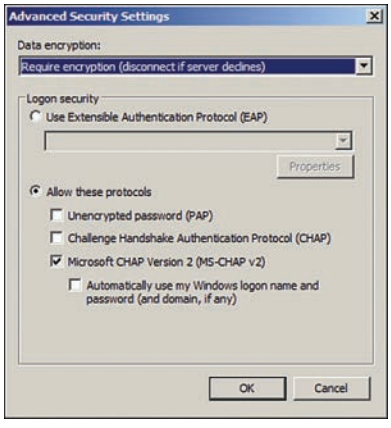


FIGURE 24-8 MS-CHAP Enabled on a Dial-Up Connection

NOTE You should use CHAP, MS-CHAP, or EAP for dial-up connections. Also, you should verify that it is configured properly on the RAS server and dial-up client to ensure a proper handshake.

The CHAP authentication scheme consists of several steps. It authenticates a user or network host to entities such as Internet access providers. CHAP periodically verifies the identity of the client by using a three-way handshake. The verification is based on a shared secret. After the link has been established, the authenticator sends a challenge message to the peer. The encrypted results are compared, and finally the client is either authorized or denied access.

The actual data transmitted in these RAS connections is encrypted as well. By default, Microsoft RAS connections are encrypted by the RSA RC4 algorithm.

You might wonder why dial-up connections are still relevant. Well, they are important for two reasons. First, the supporting protocols, authentication types, and encryption types are used in other technologies; this is the basis for those systems. Second, as previously mentioned, some organizations still use dial-up connections—for remote users or for administrative purposes. And hey, don't downplay the dial-up connection. Old-school dial-up guys used to tweak the connection to the point

where it was as fast as some DSL versions and as reliable. So, there are going to be die-hards out there as well. Plus, some areas of the United States and the rest of the world have no other option than dial-up.

However, RAS now has morphed into something that goes beyond just dial-up. VPN connections that use dial-up, cable Internet, fiber, and so on are all considered remote access.

RADIUS versus TACACS+

The ***Remote Authentication Dial-In User Service (RADIUS)*** can be used in combination with a small office/home office (SOHO) router to provide strong authentication. RADIUS provides centralized administration of dial-up, VPN, and wireless authentication and can be used with EAP and 802.1X. To set this up on a Windows server, you must load the Internet Authentication Service; it is usually set up on a separate physical server. RADIUS is a client/server protocol that runs on the application layer of the OSI model.

RADIUS works within the AAA concept: it is used to authenticate users, authorize them to services, and account for the usage of those services. RADIUS checks whether the correct authentication scheme such as CHAP or EAP is used when clients attempt to connect. It commonly uses port 1812 for authentication messages and port 1813 for accounting messages (both of which use UDP as the transport mechanism). In some proprietary cases, it uses ports 1645 and 1646 for these messages, respectively. You should memorize these four ports for the exam!

Another concept you will encounter is that of *RADIUS federation*. Here, an organization has multiple RADIUS servers—possibly on different networks—that need to communicate with each other in a safe way. This communication is accomplished by creating trust relationships and developing a core to manage those relationships as well as the routing of authentication requests. It is often implemented in conjunction with 802.1X. This federated network authentication could also span between multiple organizations.

NOTE Another protocol similar to RADIUS—though not as commonly used—is the Diameter protocol. Once again, it's an AAA protocol. It evolves from RADIUS by supporting TCP or SCTP, but not UDP. It uses port 3868.

The Terminal Access Controller Access Control System (TACACS) is another remote authentication protocol that was used more often in UNIX networks. In UNIX, the TACACS service is known as the TACACS daemon. The newer and more commonly used implementation of TACACS is called ***Terminal Access Controller Access Control System Plus (TACACS+)***. It is not backward compatible with TACACS, however. TACACS+ and its predecessor, XTACACS, were developed

by Cisco. TACACS+ uses inbound port 49 like its forerunners; however, it uses TCP as the transport mechanism instead of UDP. Let's clarify: the older TACACS and XTACACS technologies are not commonly seen anymore. The two common protocols for remote authentication used today are RADIUS and TACACS+.

There are a few differences between RADIUS and TACACS+. Whereas RADIUS uses UDP as its transport layer protocol, TACACS+ uses TCP as its transport layer protocol, which is usually seen as a more reliable transport protocol (though each has its own unique set of advantages). Also, RADIUS combines the authentication and authorization functions when dealing with users; however, TACACS+ separates these two functions into two separate operations that introduce another layer of security. It also separates the accounting portion of AAA into its own operation.

RADIUS encrypts only the password in the access-request packet, from the client to the server. The remainder of the packet is unencrypted. Other information such as the username can be easily captured, without need of decryption, by a third party. However, TACACS+ encrypts the entire body of the access-request packet. So, effectively TACACS+ encrypts entire client/server dialogues, whereas RADIUS does not. Finally, TACACS+ provides for more types of authentication requests than RADIUS.

Table 24-2 summarizes the local and remote authentication technologies covered thus far.

Key Topic

Table 24-2 Authentication Technologies

Authentication Type	Description
802.1X	An IEEE standard that defines port-based network access control (PNAC). 802.1X is a data link layer authentication technology used to connect devices to a LAN or WLAN. It defines EAP.
Kerberos	An authentication protocol designed at MIT that enables computers to prove their identity to each other in a secure manner. It is used most often in a client/server environment; the client and server both verify each other's identity.
CHAP	An authentication scheme used by the Point-to-Point Protocol (PPP) that is the standard for dial-up connections. It utilizes a challenge-response mechanism with one-way encryption. Derivatives include MS-CHAP and MS-CHAPv2.
RADIUS	A protocol used to provide centralized administration of dial-up, VPN, and wireless authentication. It can be used with EAP and 802.1X. It uses ports 1812 and 1813, or 1645 and 1646, over a UDP transport.
TACACS+	Remote authentication developed by Cisco, similar to RADIUS but separates authentication and authorization into two separate processes. It uses port 49 over a TCP transport.

Access Control Schemes

Key Topic

Access control models are methodologies in which admission to physical areas and, more important, computer systems is managed and organized. Access control, also known as an access policy, is extremely important when it comes to users accessing secure or confidential data. Some organizations also practice concepts such as separation of duties, job rotation, and least privilege. By combining these best practices along with an access control model, you can develop a robust plan concerning how users access confidential data and secure areas of a building.

There are several models for access control, each with its own special characteristics that you should know for the exam.

Discretionary Access Control

Discretionary access control (DAC) is an access control policy generally determined by the owner. Objects such as files and printers can be created and accessed by the owner. Also, the owner decides which users are allowed to have access to the objects, and what level of access they may have. The levels of access, or permissions, are stored in access control lists (ACLs).

Originally, DAC was described in The Orange Book as the Discretionary Security Policy and was meant to enforce a consistent set of rules governing limited access to identified individuals. The Orange Book's proper name is the *Trusted Computer System Evaluation Criteria (TCSEC)*, and was developed by the U.S. Department of Defense (DoD); however, The Orange Book is old (it's referred to in the movie *Hackers* in the 1990s!), and the standard was superseded in 2005 by an international standard called the Common Criteria for Information Technology Security Evaluation (or simply Common Criteria). But the DAC methodology lives on in many of today's personal computers and client/server networks.

NOTE An entire set of security standards known as the "Rainbow Series" was published by the DoD in the 1980s and 1990s. Although The Orange Book is the centerpiece of the series (maybe not in the color spectrum, but as far as security content), there are other ones you might come into contact with, such as The Red Book, which is the Trusted Network Interpretation standard. Some of the standards have been superseded, but they contain the basis for many of today's security procedures.

An example of DAC would be a typical Windows computer with two users. User A can log on to the computer, create a folder, stock it with data, and then finally configure permissions so that only she can access the folder. User B can log on to the computer but cannot access User A's folder by default, unless User A says so

and configures it as so! However, User B can create his own folder and lock down permissions in the same way. Let's say that there is a third user, User C, who wants both User A and User B to have limited access to a folder that he created. That is also possible by setting specific permission levels, as shown in Figure 24-9. The first Properties window shows that User C (the owner) has Full Control permissions. This permission is normal because User C created the folder. But in the second Properties window, you see that User A has limited permissions, which were set by User C.

Key Topic

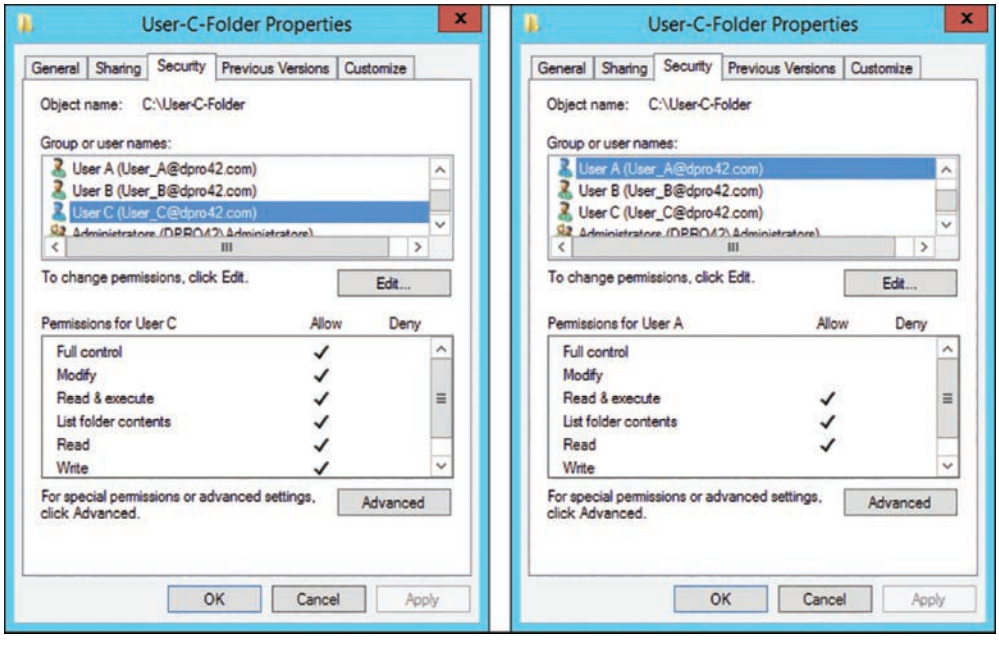


FIGURE 24-9 Discretionary Access in Windows

NOTE Take notice of standard naming conventions used in your organization. In Figure 24-9 the naming convention is *user@domainname*—for example, User_A@dpro42.com.

NOTE The owner of a resource controls the permissions to that resource! This is the core of the DAC model.

Windows networks/domains work in the same fashion. Access to objects is based on which user created them and what permissions that user assigned to those objects. However, in Windows networks, you can group users together and assign

permissions by way of roles as well. For more details, see the "Role-Based Access Control" section.

In a way, DAC, when implemented in client/server networks, is sort of a decentralized administration model. Even though you, as administrator, still have control over most, or all, resources (depending on company policy), the owners retain a certain amount of power over their own resources. But many companies take away the ability for users to configure permissions. They may create folders and save data to them, but the permissions list is often generated on a parent folder by someone else and is inherited by the subfolder.

There are two important points to remember about the DAC model. First, every object in the system has an owner, and the owner has control over its access policy. And second, access rights, or permissions, can be assigned by the owner to users to specifically control object access.

Mandatory Access Control

Mandatory access control (MAC) is an access control policy determined by a computer system, not by a user or owner, as it is in DAC. Permissions are predefined in the MAC model. Historically, it has been used in highly classified government and military multilevel systems, but you will find lesser implementations of it in today's more common operating systems as well. The MAC model defines sensitivity labels that are assigned to *subjects* (users) and *objects* (files, folders, hardware devices, network connections, and so on). A subject's label dictates its security level, or level of trust. An object's label dictates what level of clearance is needed to access it, also known as a trust level (this is also known as *data labeling*). The access controls in a MAC system are based on the security classification of the data and "need-to-know" information—where a user can access only what the system considers absolutely necessary. Also, in the MAC model, data import and export are controlled. MAC is the strictest of the access control models.

An example of MAC can be seen in FreeBSD version 5.0 and higher. In this operating system, access control modules can be installed that allow for security policies that label subjects and objects. Policies are enforced by administrators or by the OS; this is what makes it mandatory and sets it apart from DAC. Another example is Security-Enhanced Linux (SELinux), a set of kernel modifications to Linux that supports DoD-style mandatory access controls such as the requirement for a trusted computing base (TCB). Though often interpreted differently, a TCB can be described as the set of all hardware and software components critical to a system's security and all associated protection mechanisms. The mechanisms must meet a certain standard, and SELinux helps accomplish this by modifying the kernel of the Linux OS in a secure manner. Like DAC, MAC was also originally defined in

The Orange Book, but as the Mandatory Security Policy—a policy that enforces access control based on a user's clearance and by the confidentiality levels of the data.

NOTE *Rule-based* access control that uses labels is part of mandatory access control and should not be confused with *role-based* access control.

NOTE Other related access control models include Bell-LaPadula, Biba, and Clark-Wilson. Bell-LaPadula is a state machine model used for enforcing access control in government applications. It is a less common multilevel security derivative of mandatory access control. This model focuses on data confidentiality and controlled access to classified information. The Biba integrity model describes rules for the protection of data integrity. Clark-Wilson is another integrity model that provides a foundation for specifying and analyzing an integrity policy for a computing system.

Role-Based Access Control

Role-based access control (RBAC) is an access model that, like MAC, is controlled by the system, and, unlike DAC, not by the owner of a resource. However, RBAC is different from MAC in the way that permissions are configured. RBAC works with sets of permissions instead of individual permissions that are label-based. A set of permissions constitutes a role. When users are assigned to roles, they can then gain access to resources. A role might be the ability to complete a specific operation in an organization as opposed to accessing a single data file. For example, a person in a bank who wants to check a prospective client's credit score would be attempting to perform a transaction that is allowed only if that person holds the proper role. So roles are created for various job functions in an organization. Roles might have overlapping privileges and responsibilities. Also, some general operations can be completed by all the employees of an organization. Because there is overlap, an administrator can develop role hierarchies; these define roles that can contain other roles or have exclusive attributes.

Think about it. Did you ever notice that an administrator or root user is extremely powerful? Perhaps too powerful? And standard users are often not powerful enough to respond to their own needs or fix their own problems? Some operating systems counter this problem by creating mid-level accounts such as auditors (Microsoft) or operators (Solaris), but for large organizations, this approach is not flexible enough. Currently, more levels of roles and special groups of users are implemented in newer operating systems. RBAC is used in database access as well and is becoming more common in the health-care industry and government.

Attribute-Based Access Control

Attribute-based access control (ABAC) is an access model that is dynamic and context-aware. Access rights are granted to users through the use of multiple policies that can combine various user, group, and resource attributes together. It makes use of IF-THEN statements based on the user and requested resource. For example, *if* David is a system administrator, *then* allow full control access to the \\dataserver\adminfolder share. If implemented properly, this solution can be more flexible.

Rule-Based Access Control

Rule-based access control, also known as label-based access control, defines whether access should be granted or denied to objects by comparing the object label and subject label. Rule-based access control is another model that can be considered a special case of attribute-based access control (ABAC). In reality, this is not a well-defined model and includes any access control model that implements some sort of rule that governs the access to a resource. Usually, rule-based access controls are used in the context of access list implementation to access network resources—for example, where the rule is to provide access only to certain IP addresses or only at certain hours of the day. In this case, the IP addresses are attributes of the subject and object, and the time of day is part of the environment attribute evaluation.

Conditional Access

Conditional access is an access control model where access is granted based on specific criteria requirements. It is primarily used in the software as a service (SaaS) environments such as Microsoft 365 and other services provided by Microsoft Azure. The main function of conditional access is based on the concept of limiting access based on whether or not specific conditions are met. An example of those conditions might be things like IP address, type of browser, type of operating system, or geographic location.

Privileged Access Management

Privileged access management (PAM) is a system used to centrally manage access to privileged accounts. It is primarily based on the concept of least privilege. Typically, a privileged access management system is used to securely store the elevated credentials used by an organization and broker the use of those credentials based on criteria set by a privileged access management administrator. Many different privileged access management solutions are available today with varying features and functions.

Summary of Access Control Models

Table 24-3 summarizes the access control models just discussed: DAC, MAC, RBAC, ABAC, conditional access, PAM, and rule-based access control.

Key Topic

Table 24-3 Access Control Models

Access Control Model	Key Points
DAC	Every object in the system has an owner. Permissions are determined by the owner.
MAC	Permissions are determined by the system. It can be rule-based or lattice-based. Labels are used to identify security levels of subjects and objects.
RBAC	It is based on roles or sets of permissions involved in an operation. It is controlled by the system.
ABAC	It is context-aware and provides dynamic authentication. It uses IF-THEN statements to allow access.
Conditional Access	Access is granted based on specific conditions being met. Office 365 is a use case.
Privileged access management (PAM)	It provides centralized control and management of privileged credentials.
Rule-based access control	Access is granted or denied to objects by comparing the object label and subject label. Usually, it is used in the context of access list implementation to access network resources.

NOTE Another type of access control method is known as anonymous access control—for example, access to an FTP server. This method uses attributes before access is granted to an object. Authentication is usually not required.

NOTE In general, access control can be centralized or decentralized. *Centralized* access control means that one entity is responsible for administering access to resources. *Decentralized* access control means that more than one entity is responsible, and those entities are closer to the actual resources than the entity would be in a centralized access control scenario.

Access Control Wise Practices

After you decide on an access control model that fits your needs, you should consider employing some other concepts. Some of these are used in operating systems automatically to some extent:

- **Implicit deny:** This concept denies all traffic to a resource unless the users generating that traffic are specifically granted access to the resource. Even if permissions haven't been configured for the user in question, that person is still denied access. This is a default setting for access control lists on a Cisco router. It is also used by default on Microsoft computers to a certain extent. Figure 24-10 shows an example. In the folder's permissions, you can see that the Users group has the Read & Execute, List Folder Contents, and Read permissions set to Allow. But other permissions such as Modify are not configured at all—not set to Allow or Deny. Therefore, the users in the Users group cannot modify data inside the folder because that permission is implicitly denied. Likewise, they can't take full control of the folder.

NOTE The Implicit deny setting denies users access to a resource unless they are specifically allowed access.

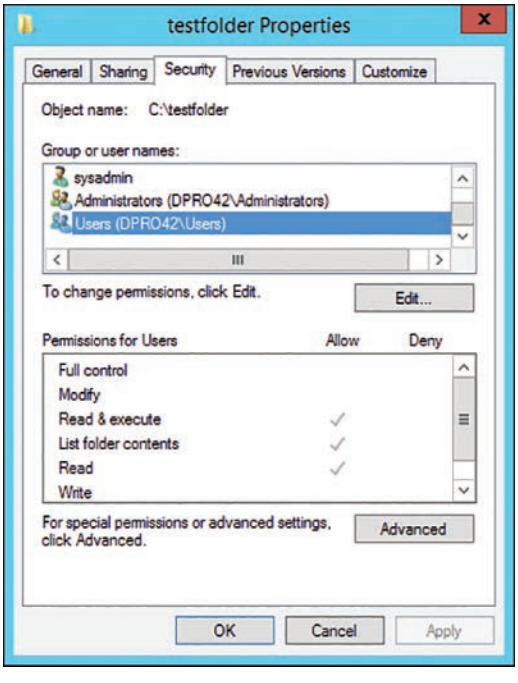


FIGURE 24-10 Implicit Deny on a Windows Folder

- **Least privilege:** With this setting, users are given only the number of privileges needed to do their job and not one iota more. A basic example would be the Guest account in a Windows computer. This account (when enabled) can surf the web and use other basic applications but cannot make any modifications to the computer system. However, least privilege as a principle goes much further. One of the ideas behind this principle is to run the user session with only the processes necessary, thus reducing the amount of CPU power needed. This hopefully leads to better system stability and system security. Have you ever noticed that many crashed systems are due to users trying to do more than they really should be allowed? Or more than the computer can handle? The concept of *least* privilege tends to be absolute, whereas an absolute solution isn't quite possible in the real world. It is difficult to gauge exactly what the "least" number of privileges and processes would be. Instead, as security administrator, you should practice the implementation of *minimal* privilege, reducing what a user has access to as much as possible. Programmers also practice this principle when developing applications and operating systems, making sure that apps have only the least privilege necessary to accomplish what they need to do. This concept is also known as "the principle of least privilege."

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 24-4 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 24-4 Key Topics for Chapter 24

Key Topic Element	Description	Page Number
Section	Authentication Management	655
Section	Authentication/Authorization	657
Section	802.1X and EAP	664
Figure 24-4	An 802.1X-enabled network adapter in Windows	665
Figure 24-5	Components of a typical 802.1X authentication procedure	666

Key Topic Element	Description	Page Number
Figure 24-6	Active Directory showing user objects	667
Figure 24-7	Results of the **netstat -an** command on a Windows server	669
Figure 24-8	MS-CHAP enabled on a dial-up connection	671
Table 24-2	Authentication Technologies	673
Paragraph	Access control schemes	674
Figure 24-9	Discretionary access in Windows	675
Table 24-3	Access Control Models	679

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

password keys, password vault, Trusted Platform Module (TPM), hardware security modules (HSMs), knowledge-based authentication (KBA), single sign-on (SSO), Security Assertion Markup Language (SAML), OAuth, OpenID, 802.1X, Extensible Authentication Protocol (EAP), Kerberos, Challenge-Handshake Authentication Protocol (CHAP), Password Authentication Protocol (PAP), Remote Authentication Dial-In User Service (RADIUS), Terminal Access Controller Access-Control System Plus (TACACS+), access control models, discretionary access control (DAC), mandatory access control (MAC), role-based access control (RBAC), attribute-based access control (ABAC), rule-based access control, conditional access, privileged access management (PAM)

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is a security model where users are given only the number of privileges needed to do their job?
2. What concept denies all traffic to a resource unless the users who generate the traffic are specifically granted access to the resource?
3. What kind of file system permissions are broken down into read, write, and execute?
4. What is an access model based on roles or sets of permissions involved in an operation?

5. What is an access model where access is controlled by the owner?
6. What is a system used to centrally manage access to privileged accounts?
7. What is an access model where permissions are determined by the system?
8. What is an authentication protocol designed by MIT that enables computers to prove their identity to each other in a secure manner?
9. What is an access control model that is dynamic and context-aware?
10. What is a physical device that can act as a secure cryptoprocessor?
11. What is an authentication based on knowledge of information associated with an individual?



This chapter covers the following topics related to Objective 3.9 (Given a scenario, implement public key infrastructure) of the CompTIA Security+ SY0-601 certification exam:

- Public key infrastructure (PKI)
  - Key management
  - Certificate authority (CA)
  - Intermediate CA
  - Registration authority (RA)
  - Certificate revocation list (CRL)
  - Certificate attributes
  - Online Certificate Status Protocol (OCSP)
  - Certificate signing request (CSR)
  - CN
  - Subject alternative name
  - Expiration
- Types of certificates
  - Wildcard
  - Subject alternative name
  - Code signing
  - Self-signed
  - Machine/computer certificate
  - Email
  - User
  - Root
  - Domain validation
  - Extended validation
- Certificate formats
  - Distinguished encoding rules (DER)
  - Privacy enhanced mail (PEM)
  - Personal information exchange (PFX)
  - .cer
  - P12
  - P7B
- Concepts
  - Online vs. offline CA
  - Stapling
  - Pinning
  - Trust model
  - Key escrow
  - Certificate chaining

Implementing Public Key Infrastructure

In this chapter we briefly dig into the vast topic of public key infrastructure (PKI). This concept is important in many aspects of security today and is essential to understand for the Security+ exam. Many of the security controls discussed in this book rely in some way on public key infrastructure.

PKI is a set of identities, roles, policies, and actions for the creation, use, management, distribution, and revocation of digital certificates. The reason that PKI exists is to enable the secure electronic transfer of information for many different purposes. You probably know that using simple passwords is an inadequate authentication method. PKI provides a more rigorous method to confirm the identity of the parties involved in the communication and to validate the information being transferred.

PKI binds public keys with the identities of people, applications, and organizations. This "binding" is maintained by the issuance and management of digital certificates by a certificate authority (CA).

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 25-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 25-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Public Key Infrastructure	1–2
Types of Certificates	3–6
Certificate Formats	7
PKI Concepts	8–9

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is a type of key used in PKI that should always be kept secret and stored securely?
   a. Private key
   b. Public key
   c. Wireless key
   d. None of these answers are correct.

2. Which of the following is utilized by a web server and viewed by the web browser?
   a. Private key
   b. Public key
   c. Web key
   d. All of these answers are correct.

3. Which standard format is used for most certificates?
   a. X.509
   b. EFS
   c. X.409
   d. PEM

4. Which of the following is a component of an X.509 certificate that includes information such as serial number and digital signature?
   a. Hash name
   b. Certificate code
   c. Bar code
   d. Certificate authority

5. Which of the following certificates allow for connections to the main website as well as subdomains?
   a. Directory certificates
   b. Public certificates
   c. Extended certificates
   d. Wildcard certificates

6. Which of the following fields in a certificate can specify additional hostnames?
   a. Hostname
   b. Subject Alternative Name
   c. Validation name
   d. None of these answers are correct.

7. Which of the following is a restricted version of BER in that it allows the use of only one encoding type?
   a. DER
   b. CER
   c. P12
   d. BER

8. Which is a way to add a layer of security to avoid root CA compromise?
   a. Online
   b. Offline
   c. Multipath
   d. Singlepath

9. Which is a method to use for blocking on-path attacks?
   a. Online
   b. Pinning
   c. Offline
   d. Chaining

Foundation Topics

Public Key Infrastructure

The following sections highlight important public key infrastructure concepts starting with key management. From there we dive into certificate authority (CA), intermediate CA, and registration authority (RA). The discussion of PKI continues with coverage of certificate revocation lists (CRLs), certificate attributes, Online Certificate Status Protocol (OCSP), and certificate signing requests (CSRs). We conclude with an overview of common name (CN), Subject Alternative Name, and expiration.

Key Management

Key Topic

Public key infrastructure (PKI) is a ***key management*** system of hardware and software, policies and procedures, and people. It is used to create, distribute, manage, store, and revoke digital certificates. If you have connected to a secure website in the past, you have utilized a PKI. However, a PKI can be used for other things as well, such as secure email transmissions and secure connections to remote computers and remote networks. The PKI is all encompassing: it includes users, client computers, servers, services, and most of all, encryption. Don't confuse PKI with public key encryption. Though they are related, PKI is a way of accomplishing public key encryption, but not all public key encryption schemes are PKI.

PKI creates asymmetric key pairs, with a public key and a private key. The private key is kept secret, whereas the public key can be distributed. If the key pair is generated at a server, it is considered to be centralized, and the public key is distributed as needed. If the key pair is generated at a local computer, it is considered to be decentralized, and the keys are not distributed; instead, they are used by that local system. An example of public key usage would be a certificate obtained by a web browser during an encrypted session with an e-commerce website. An example of private key usage would be when a user needs to encrypt the digital signature of a private email. The difference is the level of confidentiality. The public key certificate obtained by the web browser is public and might be obtained by thousands of individuals. The private key used to encrypt the email is not to be shared with anyone.

In a nutshell, public key infrastructures are set up in such a way so as to bind public keys with user identities. This is usually done with certificates distributed by a certificate authority. Less commonly, it is done by means of a web of trust.

Certificate Authorities

A ***certificate authority (CA)*** is the entity (usually a server) that issues certificates to users. In a PKI system that uses a CA, the CA is known as a trusted third party. Most PKI systems use a CA. The CA is also responsible for verifying the identity of the recipient of the certificate. An example of a technology that uses certificates would be secure websites. If you opened your browser and connected to a secure site, the browser would first check the certificate that comes from DigiCert or another similar company; it would *validate* the certificate. You (the user) and the website are the two parties attempting to communicate. The CA is a third party that negotiates the security of the connection between you and the website.

For a user to obtain a *digital identity certificate* from a CA, the user's computer must initiate a *certificate signing request (CSR)* and present two items of information: The first is proof of the user's identity; the second is a public key. This public key is then matched to the CA's private key, and if successful, the certificate is granted to the user.

A basic example of this would be if you connect to www.paypal.com. When you connect to this website, it automatically redirects you to https://www.paypal.com, which is secured by way of a DigiCert-issued certificate. You know you have been redirected to a secure site because the browser has various indicators. For instance, the web browser will probably show a padlock in the locked position, and it and the name of the company will be displayed as shown in Figure 25-1.

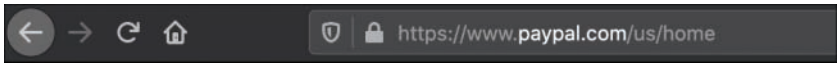


FIGURE 25-1 A Secure Connection, Shown in Firefox

If you were to click on the padlock icon area of the address field in Firefox, you could ultimately get to the certificate details. Figure 25-2 shows some of the certificate details associated with this domain name.

The figure shows when the certificate was originally issued and when it will expire, among other information. You can also note that the certificate has been fingerprinted with SHA-256 (a variant of SHA-2) enabling you or the website (or issuer) to verify the integrity of the certificate. If for some reason the certificate cannot be verified by any of the parties, and the issuer confirms this, then the issuer would need to revoke it and place it in the ***certificate revocation list (CRL)***. The Details tab gives advanced and more complete information about the certificate used. You should look at a few more websites that use SSL/TLS certificates and peruse the General and Details tabs. Compare the certificates with each other to learn more

about the different levels of encryption, different levels of fingerprinting, and the different issuing companies.

Key Topic

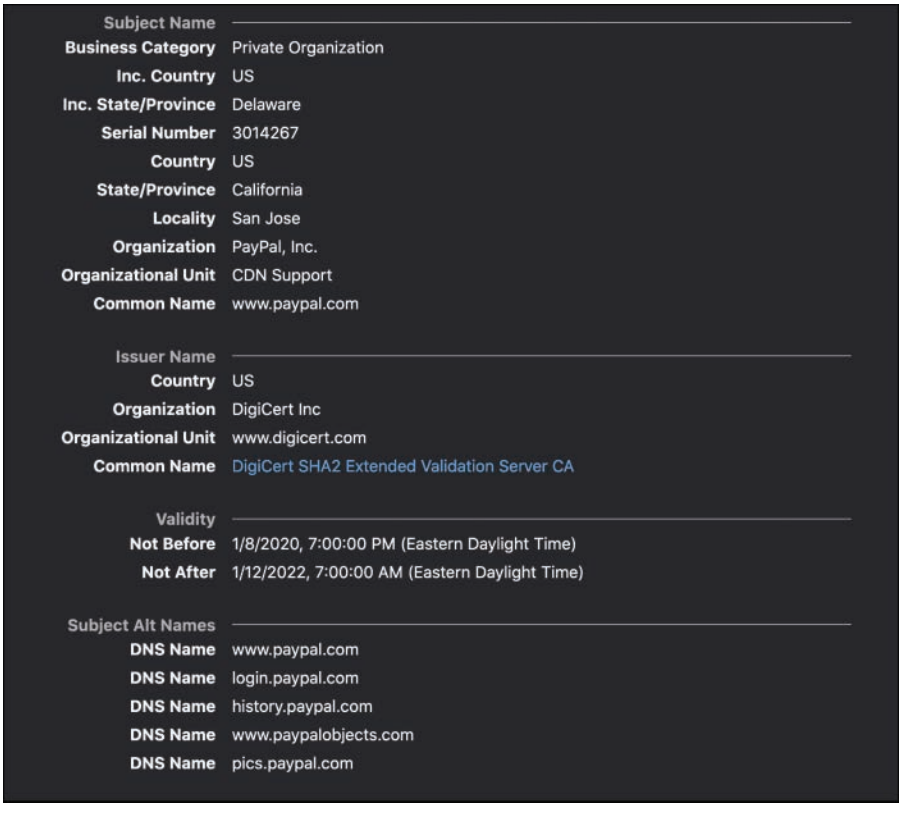


FIGURE 25-2 Details of a Typical DigiCert Certificate

Recipients can use one or more certificates. Certificate mapping defines how many certificates are associated with a particular recipient. If an individual certificate is mapped to a recipient, it is known as ***one-to-one mapping***. If multiple certificates are mapped to a recipient, it is known as ***many-to-one mapping***. Multiple certificates might be used if the recipient requires multiple secure (and separate) communications channels.

In some cases, a ***registration authority (RA)*** is used to verify requests for certificates. If the request is deemed valid, the RA informs the CA to issue the certificate. An RA might also be used if the organization deals with several CAs. In this case, the RA is at the top of a hierarchical structure and verifies the identity of the user. An RA isn't necessary in a PKI, but if you are centrally storing certificates, a CA is necessary.

Certificate authorities aren't just for the rich and famous (for example, PayPal using DigiCert as the issuer). You can have a CA, too! If you are running a Windows Server, you can install your own CA—for example, one that utilizes L2TP or possibly SSL/TLS. Of course, a server's built-in certificates are not necessarily secure. If you were to implement this technology in a secure environment in your organization, you would probably want to obtain proper certificates from a trusted source to use with the Windows Server. When implementing certificates in Windows Server, you would use the Active Directory Certificate Services (AD CS) utility. From there you can define object identifiers (OIDs), which are built into AD CS for either low, medium, or high assurance. Or, you can have Windows randomly assign them. For security purposes, obtain the OID before completing the configuration of the CA.

Certificate authorities can be subverted through the use of social engineering. If a person posing as a legitimate company managed to obtain certificates from a trusted source, those certificates would appear to be valid certificates and could cause widespread damage due to connections made by unsuspecting users—that is, until the certificates were revoked. This happens sometimes, but the CA issuer usually finds out quickly and takes steps to mitigate the problem, including revoking the certificate(s) and notifying any involved parties of the incident.

The ***certificate revocation list (CRL)*** is a list of certificates that are no longer valid or that have been revoked by the issuer. There are two possible states of revocation: revoked, which is when a certificate has been irreversibly revoked and cannot be used again, and hold, which is used to temporarily invalidate a certificate. Reasons for revoking a certificate include the compromising or theft of a certificate or entire CA, unspecified certificates, superseded certificates, held certificates, and key or encryption compromise. The CRL is published periodically, usually every 24 hours. This list enables users of an issuer's certificates to find out whether a certificate is valid. CRLs, like the certificates themselves, carry digital signatures to prevent denial-of-service and spoofing attacks; the CRL is digitally signed by the CA.

An alternative to the CRL is the ***Online Certificate Status Protocol (OCSP)***. It contains less information than a CRL does, and the client side of the communication is less complex. However, OCSP does not require encryption, making it less secure than CRL.

Certificate Attributes

The attributes in a certificate are essentially the various fields that define things like who issued the certificate and whom it is issued to. Additionally, these attributes would include information about what the certificate use intended, when it was issued, and when it will expire. These are just a few of the various attributes used in digital certificates. These attributes can then be used for authentication and

validation purposes. Table 25-2 summarizes some of the most used attributes and what they are used for.

Key Topic

Table 25-2 Certificate Attributes

Attribute	Description
Common name (CN)	The common name is the fully qualified domain name (FQDN) of the entity that the certificate is issued to. It is a field that is very often submitted incorrectly with certificate signing requests. The CN should be the same FQDN as the actual DNS name you are using in the web address you are using to access the site. If it is different from what is in the actual certificate, you will receive an error.
Organization (O)	This is the legal name of the organization that owns the site that the certificate will be used on.
Locality (L)	The locality field is used to specify the city where the legal organization is located. This should always be the full name such as North Carolina instead of NC.
Organizational unit (OU)	The OU is typically the department within the organization that will be utilizing the certificate.
Country name (C)	The country name is simply the two-letter country code for which the legal organization is located. For instance, US would be used for United States.
Serial number	This is the number issued and tracked by the CA that issued the certificate.
Issuer	This is the CA that issued this certificate. (Even root certificates need to have their certificates issued from someone, perhaps even themselves.)
Validity dates	These dates are shown in the time window during which the certificate is considered valid. If a local computer believes the date to be off by a few years, that same PC may consider the certificate invalid due to its own error about the time. Using the Network Time Protocol (NTP) is a good idea to avoid this problem.
Public key	The contents of the public key and the length of the key are often shown. After all, the public key is public.
Thumbprint algorithm and thumbprint	This is the hash for the certificate. On a new root certificate, you could use a phone to call and ask for the hash value and compare it to the hash value you see on the certificate. If it matches, you have just performed out-of-band verification (using the telephone) of the digital certificate.

Subject Alternative Name

The Subject Alternative Name is a field (or fields) in PKI certificates that allows an organization to specify additional hostnames, domain names, email addresses, or URIs for use with a single certificate. The idea of utilizing a Subject Alternative Name field in a certificate is to provide flexibility to system administrators. Figure 25-3 shows a Subject Alternative Name used for a Gmail server. Two DNS names are listed here: mail.google.com and inbox.google.com. The single certificate therefore can be used for both of these domain names.

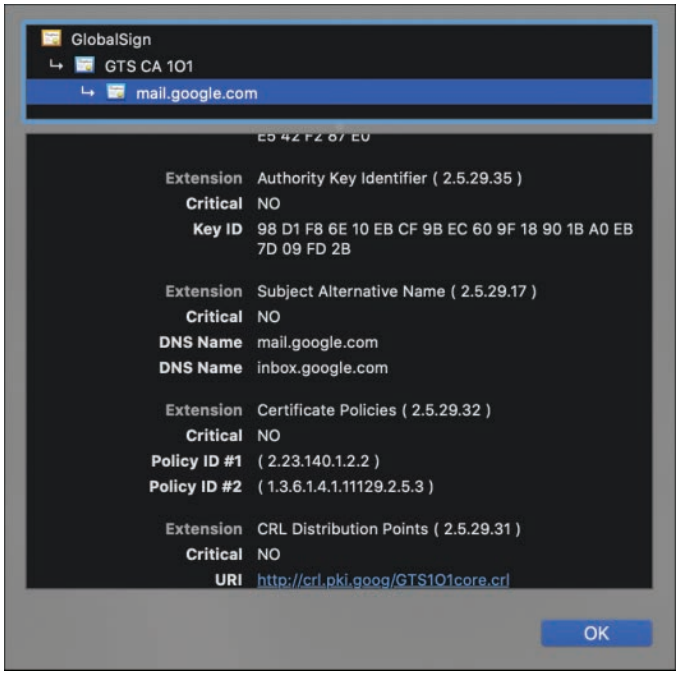


FIGURE 25-3 Subject Alternative Name

Expiration

A certificate is issued with a valid *from* and valid *to* date. This is the period of time that the certificate should be considered valid. If the certificate is encountered any time before or after those dates, the application that is validating the certificate (such as a web browser) should produce an error message stating that the certificate is not valid. It is important when issuing a certificate to issue it for a sufficient period of time. For instance, if the certificate is meant for long-term use and the application or device utilizing the certificate is not easily updated, you might want to extend the certificate validation period longer. This way, the certificate will not unexpectedly expire and break the application.

Types of Certificates

Certificates are digitally signed electronic documents that bind a public key with a user identity. The identity information might include a person's name and organization, or other details relevant to the user to whom the certificate is to be issued. Most certificates are based on the ***X.509*** standard, which is a common PKI standard developed by the ITU-T that often incorporates the single sign-on (SSO) authentication method. This way, a recipient of a single X.509 certificate has access to multiple resources, possibly in multiple locations. Although difficult, X.509 certificates that use MD5 and SHA1 hashes can be compromised. A more powerful hashing algorithm such as SHA2 should be implemented with the certificate. X.509 is the core of the Public Key Infrastructure Exchange (PKIX), which is the IETF's Public Key Infrastructure (X.509) working group. Components of an X.509 certificate include the following:

- Owner (user) information, including public key
- Certificate authority information, including name, digital signature, serial number, issue and ***expiration*** dates, and version

Certificates can be used for connections to websites, for email, and for many other things in the Internet world, as well as encryption done locally. For example, a user working in a Windows environment might want to use the Encrypting File System (EFS) to encrypt data locally. The Windows domain can be configured to allow for user certificates governing and enhancing this encryption process. So, certificates can be used internally or externally, but most people are more familiar with the certificate used to make secure HTTP connections, usually with SSL/TLS-based certificates. We focus mostly on that type of certificate as we move forward.

SSL Certificate Types

It's a good idea to classify certificates the way companies that sell them do. This includes domain, organizational, and extended validation certificates. In ***domain validation (DV)*** certificates, the certificate authority checks the rights of the applicant to use a specific domain name. Organizational validation (OV) certificates go beyond this by also conducting some vetting of the organization involved, the result of which is displayed to customers. ***Extended validation (EV)*** certificates go further by conducting a thorough vetting of the organization. Issuance of these certificates is strictly defined.

Many companies have subdomains for their websites. For example, consider h4cker.org. I might also opt to create subdomains such as tools.h4cker.org and software.h4cker.org. Generally, if you connect to a secure website that uses subdomains,

a single certificate allows for connections to the main website and the subdomains. This is known as a ***wildcard certificate***; for example, *. h4cker.org, meaning all subdomains of h4cker.org. Your organization might allow this or, for additional security, might use a different certificate for each subdomain (possibly from different providers), but this approach can prove to be expensive. For small businesses and organizations, the single certificate is usually enough. In fact, if the provider allows it, a small organization can use a multidomain certificate. By modifying the ***Subject Alternative Name (SAN)*** field, an organization can specify additional hostnames, domain names, IP addresses, and so on. Table 25-3 provides a summary of the various certificate types and how they are used.

Key Topic

Table 25-3 Certificate Types

Attribute	Description
Wildcard	Many companies have subdomains for their websites. For example, for h4cker.org, I might also opt to create subdomains such as tools. h4cker.org and software.h4cker.org. Generally, if you connect to a secure website that uses subdomains, a single certificate will allow for connections to the main website and the subdomains. This is known as a wildcard certificate; for example, *. h4cker.org, meaning all subdomains of h4cker.org. Your organization might allow this or, for additional security, might use a different certificate for each subdomain (possibly from different providers), but this can prove to be expensive. For small businesses and organizations, the single certificate is usually enough.
Subject Alternative Name	The Subject Alternative Name is a field (or fields) in PKI certificates that allows an organization to specify additional hostnames, domain names, email addresses, or URIs for use with a single certificate. The purpose of utilizing a Subject Alternative Name field in a certificate is to provide flexibility to system administrators.
Code signing	A code signing certificate is typically used by software developers. They utilize it to sign the code that they have created. This can be in the form of a software driver, an application, or an executable program. The intent is to provide the end user with the ability to verify that the code has not been tampered with.
Self-signed	A self-signed certificate is created and used on a system that not only creates the certificate but also signs it. The system signs it with its own private key. This type of certificate does not offer the authentication usage that a publicly signed certificate would. When you encounter a self-signed certificate with your web browser, you will receive an error that the certificate is not valid because it is not signed by a trusted certificate authority.

Attribute	Description
Machine/computer	A machine certificate is also known as a computer certificate. A computer certificate is typically used when authenticating a computer or user connecting to a network via a VPN or 802.1x. With a user certificate, the Subject Alternative Name field in the certificate contains the fully qualified domain name, which is used in the authentication process.
Email	An email certificate is meant to be used for signing and encrypting email communications.
User	A user certificate is typically used when authenticating a computer or user connecting to a network via a VPN or 802.1x. With a user certificate, the Subject Alternative Name field in the certificate contains the User Principal Name (UPN), which is used in the authentication process.
Root	A root certificate contains the public key of the CA server and the other details about the CA server.
Domain validation	In DV certificates, the certificate authority checks the rights of the applicant to use a specific domain name.
Extended validation	EV certificates conduct a thorough vetting of the organization. Issuance of these certificates is strictly defined.

Certificate Chaining

A certificate chain is a list of certificates that utilizes the chain of trust concept. Each component of the system is validated from the bottom up. In a certificate chain, also known as a certification path, the anchor for this trust is the ***root certificate authority***. At the bottom of the chain is the end-entity certificate for a ***machine/computer certificate***. That's the certificate that you would see if you looked at the details of a secure HTTPS session in your web browser. It then handshakes with an intermediate certificate belonging to an ***intermediate certificate authority*** (CA). This certificate signs the end-entity certificate. It then handshakes with the root certificate, which represents the ***root certificate authority***. It signs the ***intermediate certificate*** and in and of itself is self-signing, which means that it not only creates the certificate but also signs it with its own private key. The root CA employs ***code signing***: digitally signing and timestamping the certificate to provide integrity and authenticity. However, even this can be defeated, so as security administrator, you always have to be on the lookout for CVEs detailing revoked certificates and even entire issuing certificate companies that may have been compromised.

Certificate Formats

Key Topic

You should know several certificate formats for the Security+ exam. They can be identified in part by their file extension or encoding type used. First, let's look at the ITU-T X.690 encoding formats:

- **Basic Encoding Rules (BER):** This is the original ruleset governing the encoding of the ASN.1 data structure. Any data created is encoded with a type identifier, a length description, and the content's value. BER can use one of several encoding methods.
- **Canonical Encoding Rules (CER):** This is a restricted version of BER in that it allows the use of only one encoding type; all others are restricted.
- ***Distinguished Encoding Rules (DER)*:** Another restricted variant of BER, this allows for only one type of encoding and has restrictive rules for length, character strings, and how elements are sorted. It is widely used for X.509 certificates. For example, certificate enrollment in Windows Servers uses DER exclusively.

Now, let's briefly examine the certificate formats and extensions you might encounter. ***PEM*** is a common format that uses base64-encoded ASCII files. It stands for Privacy-enhanced Electronic Mail and can be identified with the *.pem* file extension, though the format might also use .crt (for example, Microsoft), .cer, or .key extensions. This format uses the DER encoding method. If the certificate uses a file extension different from .pem, you can tell whether it is a PEM by opening the file with a text editor and looking for the Begin Certificate and End Certificate statements. However, if the certificate uses the .der extension, then the certificate file is in binary form instead of ASCII. Because it is in binary, you will not see the Begin and End Certificate statements that are displayed in a .pem.

P12/PFX is a binary format based on PKCS#12 used to store a server certificate, intermediate certificates, and the private key in one encryptable file. It is typically used to import and export certificates and private keys. You may see the *.pfx* and *.p12* extensions associated with PKCS#12-based files. *.pfx* stands for Personal Information Exchange and is used by Microsoft for release signing. The certificate and its private and public keys are stored in the *.pfx* file. A *.pfx* file can also be developed by combining a private key with a PKCS #7 ***.p7b*** file, as might be done in Windows Internet Information Services (IIS). Or, *.p7b* format certificates can be used by themselves in IIS as the basis for S/MIME and single sign-on.

NOTE Some of these extensions are also used for different types of data such as private keys, and not only for certificates. It is also possible to convert from one format to another using tools such as OpenSSL.

Key Topic PKI Concepts

The following sections discuss PKI concepts such as online vs. offline CA, certificate stapling, pinning, trust model, key escrow, and certificate chaining.

Trust Model

A ***web of trust*** is a ***decentralized trust model*** that addresses issues associated with the public authentication of public keys common to CA-based PKIs. It is considered peer-to-peer in that there is no root CA; instead, ***self-signed certificates*** are created and used that have been attested to by the creator. Users can decide what certificates they want to trust and can share those trusted certificates with others, making the web of trust grow larger. Of course, one of the most common reasons that a certificate issuer is not recognized by a web browser is due to unknown self-signed certificates. This model can also interoperate with standard CA architectures inherent to PKI. The more people who show trust of a certificate, the higher the chance that it is legitimate. This model is used by PGP, which enables users to start their own web of trust, self-publishing their own public key information.

Certificate Pinning

Key Topic

One way to add security to the certificate validation process is to use certificate ***pinning***, also known as SSL pinning or public key pinning. It can help detect and block many types of on-path attacks by adding an extra step beyond normal X.509 certificate validation. Essentially, a client obtains a certificate from a CA in the normal way but also checks the public key in the server's certificate against a hashed public key used for the server name. This functionality must be incorporated into the client side, so it is important to use a secure and up-to-date web browser on each client in order to take advantage of certificate pinning.

Stapling, Key Escrow, Certificate Chaining, Online vs. Offline CA

An alternative to OCSP is OCSP ***stapling*** (previously known as TLS Certificate Status Request), which allows the presenter of the certificate to bear the cost involved when providing OCSP responses.

Certificate keys can also be held in escrow. In ***key escrow*** a secure copy of a user's private key is held in case the key is lost. This may be necessary so that third parties such as government or other organizations can ultimately gain access to communications and data encrypted with that key. If data loss is unacceptable, you should implement key escrow in your PKI.

When installing a certificate authority to a Windows Server, you can set up a recovery agent for lost or corrupted keys. To do this, you need Windows Server and need to set up an enterprise-level CA. In this configuration, the certificates (or private keys) are archived at the CA. If a ***key recovery agent*** has been configured, lost, or corrupted, keys can be restored. It's important to use some type of software that can archive and restore keys in case of an incident or disaster.

Another way to avoid single points of failure, such as a single CA, is to organize certificate authorities in a hierarchical manner. At the top of the tree is a root CA; underneath are subordinate, or intermediate, CAs that offer redundancy and can sign certificates on behalf of the root CA. Though CA exclusivity is common, it is not the only type of architecture used to bind public keys to users. In some cases, a centralized model for certificates is not required or desired.

If a root CA is compromised, all of its certificates are then also compromised, which could affect an entire organization and beyond. The entire certificate chain of trust can be affected. One way to add a layer of security to avoid root CA compromise is to set up an offline root CA. Because it is offline, it will not be able to communicate over the network with the subordinate CAs, or any other computers for that matter. Certificates are transported to the subordinate CAs *physically* using USB flash drives or other removable media. Of course, you would need to have secure policies regarding the use and transport of media and would need to incorporate data loss prevention (DLP), among other things. But the offline root CA has some obvious security advantages compared to an online root CA. You should consider this *offline* mindset when dealing with critical data and encryption methods.

NOTE One thing to take away from this discussion of certificates is that there have been many certificate exploits in the past and lots of vulnerabilities still exist. Be very careful during the planning stage of certificates.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 25-4 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 25-4 Key Topics for Chapter 25

Key Topic Element	Description	Page Number
Paragraph	Key management	688
Figure 25-2	DigiCert certificate details	690
Table 25-2	Certificate Attributes	692
Table 25-3	Certificate Types	695
Paragraph	Certificate formats	697
Section	PKI Concepts	698
Paragraph	Certificate pinning	698

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

public key infrastructure (PKI), key management, registration authority (RA), certificate revocation list (CRL), Online Certificate Status Protocol (OCSP), certificates, X.509, certificate authority, expiration, domain validation (DV), extended validation (EV), wildcard certificate, Subject Alternative Name (SAN), root certificate authority, machine/computer certificate, intermediate certificate authority, code signing, Distinguished Encoding Rules (DER), PEM, P12/PFX, .p7b, web of trust, self-signed certificates, pinning, stapling, key escrow, key recovery agent

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. In X.509, the owner does not use a _____ key.
2. What two items are included in a digital certificate?
3. Rick has a local computer that uses software to generate and store key pairs. What type of PKI implementation is this?
4. What ensures that a CRL is authentic and has not been modified?
5. What encryption concept is PKI based on?
6. You are in charge of PKI certificates. What should you implement so that stolen certificates cannot be used?
7. What should you publish a compromised certificate to?
8. You have been asked to set up authentication through PKI and encryption of a database using a different cryptographic process to decrease latency. What encryption types should you use?
9. Describe key escrow.
10. When a user's web browser communicates with a CA, what PKI element does the CA require from the browser?



This chapter covers the following topics related to Objective 4.1 (Given a scenario, use the appropriate tool to assess organizational security) of the CompTIA Security+ SY0-601 certification exam:

- Network reconnaissance and discovery
  - tracert/traceroute
  - nslookup/dig
  - ipconfig/ifconfig
  - nmap
  - ping/pathping
  - hping
  - netstat
  - netcat
  - ip scanners
  - arp
  - route
  - curl
  - theHarvester
  - sn1per
  - scanless
  - dnsenum
  - Nessus
  - Cuckoo
- File manipulation
  - head
  - tail
  - cat
  - grep
  - chmod
  - logger
- Shell and script environments
  - SSH
  - PowerShell
  - Python
  - OpenSSL
- Packet capture and replay
  - Tcpreplay
  - Tcpdump
  - Wireshark
- Forensics
  - dd
  - Memdump
  - WinHex
  - FTK imager
  - Autopsy
- Exploitation frameworks
- Password crackers
- Data sanitization

Using the Appropriate Tool to Assess Organizational Security

Security teams are constantly under attack from cybercriminals and threat actors, so they need to be able to use a mixture of different security tools to identify attacks quickly, before those attacks have a chance to cause grave damage to the business. Information security threats continually evolve, and defenses against them must evolve as well. Many best-practice frameworks, standards, and laws require a flexible response based on regular risk assessments. Various assessment tools are key to ensuring an organization is prepared and protected.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 26-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 26-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Network Reconnaissance and Discovery	1–2
File Manipulation	3–4
Shell and Script Environments	5–6
Packet Capture and Replay	7–8
Forensics	9
Exploitation Frameworks	10
Password Crackers	11
Data Sanitization	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. How can you use traceroute to shore up cybersecurity in an organization?
   a. Discovering where hosts and systems are located on a Linux-based network
   b. Tracing traffic to the Internet and where it came from
   c. Tracing all routes to Windows hosts that are running on a private network
   d. Tracing attackers that enter your network and identifying them

2. What process does nslookup use to obtain domain names from IP addresses?
   a. nslook, where a lookup is made of the NS database
   b. Reverse DNS lookup
   c. Forward DNS lookup
   d. iptrace domain lookup service

3. The command-line utility **head** can read a number of files. How many lines can it read by default?
   a. 100 lines
   b. 500 lines
   c. 10 lines
   d. 50 lines

4. The command-line tool **cat** is a universal Linux tool that allows you to view almost any file. What is the process it uses to do this?
   a. Copying the file to special space on the swapfile
   b. Copying standard input to standard output
   c. Duplicating the file to a temp file and showing the original
   d. Copying the file from the file system and displaying it to tty1

5. A shell can have two types of variables. What are they?
   a. /var/log
   b. Environment and local variables
   c. /var/variables
   d. Local variables and system variables

6. Every time it starts a session, the shell builds what area that will contain variables that define system properties?
   a. env and tmp variables
   b. Shellconf spaces
   c. The tty1 and tty2 terminals
   d. Environment

7. The Wireshark GUI tool is a packet analysis tool and is available under what platforms?
   a. CubeOS
   b. CP/M
   c. Windows, Linux, and macOS
   d. MSDos5.0

8. The packet replay tool **tcpreplay** can replay packets at different intervals. What command-line code is required?
   a. **unique ip**
   b. **pps-packets**
   c. **mbps-string**
   d. **pps-multi**

9. The disk duplication tool **dd** can be used to create a complete image of a hard disk. What is the proper syntax?
   a. **# dd hdadisk.img**
   b. **# dd if = /dev/hda of = ~/hdadisk.img**
   c. **# dd if = /dev/hda image**
   d. **# dd /dev/hda /hdadisk.img**

10. What is one of the most notable and widely used exploitation frameworks?
    a. backtrack
    b. chrona
    c. twocows
    d. Metasploit

11. Passwords are only as good as the complexity and length that the creator used. What popular password cracker enables highly parallelized password-cracking capabilities and is available on nearly every operating system platform?
    a. Rainbow Crack
    b. Active Directory Fast Crack
    c. Hashcat
    d. L0phtCrack

12. Data sanitization is the process of deliberately, permanently, and irreversibly removing or destroying the data stored on a memory device to make it unrecoverable. Which of the following is a method to ensure no data is recoverable no matter what software recovery process is used?
    a. Low-level format or fdisk /s
    b. Data erasure writing 0s and 1s
    c. Physical destruction shredding
    d. Kali Linux installed on the drive

Foundation Topics

Network Reconnaissance and Discovery

Network reconnaissance is a term for testing for potential vulnerabilities in a computer network. This may be legitimate activity by the network owner/operator who is seeking to protect it or to enforce its acceptable use policy. It also may be a precursor to external attacks on the network. Certain apparent reconnaissance activities, which would be highly suspicious if coming from outside the network, may be perfectly normal network performance and reliability monitoring when performed inside the boundaries of the network. Some network intrusion detection systems (NIDS) have difficulty in determining if a reconnaissance activity is internal or external and thus generate many false alarms causing fear, uncertainty, and doubt. The process of network reconnaissance and discovery involves running a number of tools, from address sweeps to port scanning.



tracert/traceroute

Traceroute (***tracert/traceroute***), depending on your operating system, dictates which command you use. In computing, traceroute and tracert are computer network diagnostic commands for displaying possible routes and measuring transit delays of packets across a network, including the Internet. For example if you want to see how far a particular host like google.com is located from the network you are on, you can open a command prompt and then enter **tracert www.google.com**.

To traceroute from your Windows computer, you must allow access via the Windows firewall and quite possibly your hardware firewall. When that process is complete, you can traceroute as shown in Example 26-1. As you can see, the **tracert** tool provides the number of hops, the time taken to traverse each device, and the specific host IP address along the path taken to get to the destination.

Example 26-1 tracert Output

```
C:>tracert -4 www.google.com
Tracing route to www.google.com [91.245.214.163] over a maximum of 30
hops:
1        <1 ms     < 1ms       < 1ms     xopret.usnet.net [1.2.3.10]
2        <1 ms     <1 ms       <1 ms     83.45.12.120
3        < 1 ms     <3 ms      <3 ms     83.45.100.12
4        < 1 ms     <3 ms      <3 ms     83.100.0.1
5        *           *          *         request timed out.
6        < 1 ms     <3 ms      <3 ms     secr-sarnet-bgp.secr.bgp.ac.rs
7        < 1 ms     <3 ms      <3 ms     eft-tcr-bgp.eft.bg.ac.rs
8        < 1 ms     <3 ms      3 ms      193.105.163.12
9        < 148 ms   <144 ms    <143 ms   cache.google.com [91.245.214.63]
```

Mileage may vary, meaning your stops (routers) along the path may vary depending on where in the world you are located. You may have similar results by running a Linux host under a Windows virtual machine (VM). The underlying mechanisms either work or not based on your host bypass settings. On Linux hosts, there are two options for traceroute: **traceroute** and **inetutils-traceroute**.

Figure 26-1 shows a Linux host performing a **traceroute** to google the Domain Name System (DNS).

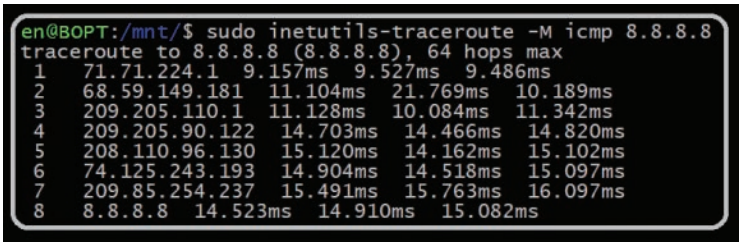


FIGURE 26-1 Kali Linux Host Performing a **traceroute** to Google DNS

NOTE Using the Linux **traceroute** (**tracert** in Windows) command allows you to document hosts' locations on your local network and map out the current location/configuration and connected devices. A baseline network diagram should be used and continuously updated to document all systems.

NOTE The **tracert** utility traces the route that a packet takes and records the hops along the way. It is a great tool to find out where a packet is getting hung up. The Trace Route utility works on Windows-based systems; **traceroute** works on macOS and Linux-based systems.

Key Topic

nslookup/dig

nslookup is a simple but practical command-line tool that is principally used to find the IP address that corresponds to a host or the domain name that corresponds to an IP address (a process called reverse DNS lookup). DNS is a major component of **nslookup**; it requires properly configured DNS. nslookup allows itself to be used in the command line of the operating system in question. Windows users start the service via the command prompt, and Linux users start it via the terminal window. Additionally, a number of services now make it possible to use nslookup online.

nslookup retrieves the relevant address information directly from the DNS cache of name servers, a process that can be achieved through two different modes that the user can choose from. In the noninteractive mode, the tool inspects the resource records (which is what the address entries in the DNS cache are called) that are stored in the local name server in a standard way. This mode is especially well suited for simple queries, for which a single domain entry needs to be looked up. When you want to use a different DNS server for the research and complete more complex search processes, you need to use interactive mode. You must start the command-line program separately in this mode at first.

The use of nslookup is appropriate in different scenarios. In this way, the command-line program is an especially crucial tool when resolving DNS issues. First and foremost, a data query helps find the cause of the problem at hand. For example, if you use Windows and experience an issue with Active Directory, you can use nslookup to quickly and easily check whether all involved servers are converted in the Domain Name System as planned. Generally, this program shows its strengths when various subdomains are involved—not only when it is simply a matter of checking them, but also when there are connection problems.

The ***dig*** command enables you to query DNS nameservers for information about host addresses, mail exchanges, nameservers, and related information. This tool can be used from any Linux or macOS operating system. The most typical use of **dig** is to simply query a single host.

Using **dig**, you can perform any valid DNS query, the most common of which are A (the IP address), TXT (text annotations), MX (mail exchanges), and NS (nameservers).

NOTE nslookup queries the DNS server to check whether the correct information is in the zone database. It is used on Windows-based systems, whereas dig, which stands for Domain Information Groper, is used on macOS and Linux-based systems.

Key Topic

ipconfig/ifconfig

ipconfig is a Windows-based operating system command that displays all current TCP/IP network configuration values and refreshes Dynamic Host Configuration Protocol and Domain Name System settings. A number of switches can be used with the command. The correct syntax is **ipconfig** */parameter_name*. Table 26-2 highlights some of the options available.

Table 26-2 ipconfig Options

Command Option	Description
ipconfig /all	Displays the full TCP/IP configuration for all adapters
ipconfig /displaydns	Displays the contents of the DNS client resolver cache
ipconfig /flushdns	Flushes and resets the contents of the DNS client resolver cache
ipconfig /release	Sends the DHCPRELEASE message to the DHCP server to release the current DHCP configuration
ipconfig /renew	Renews the DHCP configuration for all adapters
ipconfig /?	Displays the help options at the command prompt

In macOS, the **ipconfig** utility is simply a wrapper for the IPConfiguration agent. It serves to control both DHCP and BootP right from the command line. It can also be used to verify your DNS/DHCP server settings are appropriate and that the server is handing out information.

On Linux operating systems, **ifconfig** works similarly to the way **ipconfig** works on Windows systems. This command-line interface tool is also used in the system startup scripts of many operating systems. It has features for configuring, controlling, and querying TCP/IP network interface parameters.

The ***ifconfig*** (interface configuration) command is used to configure the kernel-resident network interfaces. It is used at boot time to set up the interfaces as necessary. After that, it is usually used when needed during debugging or when you need system tuning. Also, this command is used to assign the IP address and netmask to an interface or to enable or disable a given interface. The syntax for this command is as follows:

```
ifconfig [...OPTIONS] [INTERFACE]
```

Some newer Linux version distributions don't have the **ifconfig** command pre-installed. So, if you receive the error "ifconfig: command not found," you need to install it (see your distribution for the command). On Debian, the command was replaced with **ip addr**, but you can still install net-tools to obtain the **ifconfig** command.

Figure 26-2 shows a user on a Linux host checking to see what IP addresses are configured on the system. As you can see, the address shows up under inet—in this case, 192.168.43.7

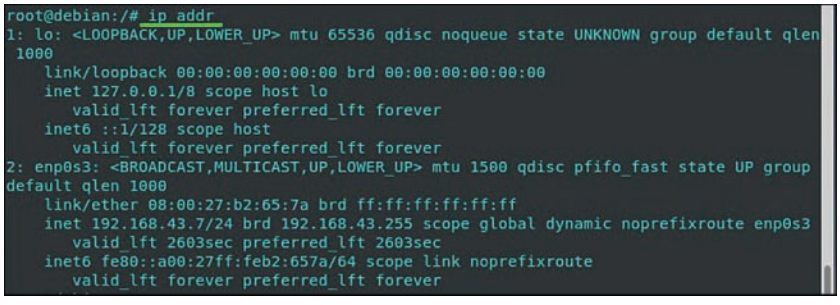


FIGURE 26-2 A Local Host IP Address Discovery

Key Topic

nmap

nmap is a popular open source vulnerability scanner. Although it was once a Linux-only tool, there was a Windows version released in 2000 and has since become the second most popular nmap platform (behind Linux). Binaries are available for nmap on newer Windows client systems, as well as Windows Server versions. Although it has improved dramatically, the Windows port is not quite as efficient as on Linux version.

nmap supports only Ethernet interfaces (including most 802.11 wireless cards and many VPN clients) for raw packet scans.

Scan speeds on Windows are generally comparable to those on Linux, although Linux often has a slight performance edge. One exception is the connect scan (-sT), which is often much slower on Windows because of deficiencies in the Windows networking API. Connect scan performance can be improved substantially by applying the registry changes in the nmap_performance.reg file included with nmap. By default, these changes are applied for you by the nmap executable installer. Most people simply check the box to apply these changes in the executable nmap installer, but you can also apply them by double-clicking on nmap_performance.reg.

Figure 26-3 shows a Windows command prompt executing an **nmap** scan of scanme.nmap.org using switches with the following options:

- **-sVC**, where -sV is the service/version detection probe to open ports to determine service/version info, and C is a script scan.
- **-O** enables operating system (OS) detection.
- **-T4** is the timing template; setting it to 4 is high speed (5 is max).

The results of this scan show that port 22 is open and this host is an Ubuntu-type operating system.

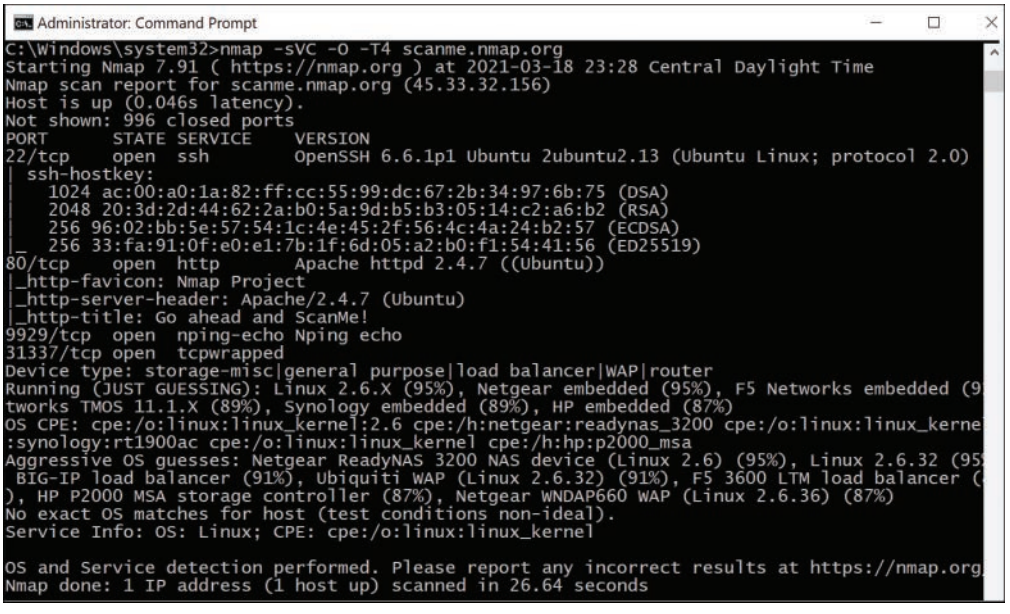


FIGURE 26-3 Windows **nmap** Scan of a Host

There are hundreds of choices when it comes to **nmap** Switch - Options. To see all of them, just enter **nmap** at the command prompt.

The Windows version of **nmap** now comes bundled with the NMAP-Zenmap GUI, which has a handy drop-down menu for the most common types of scans, or you can enter in the command area specifically what you want to execute, as demonstrated in Figure 26-4.

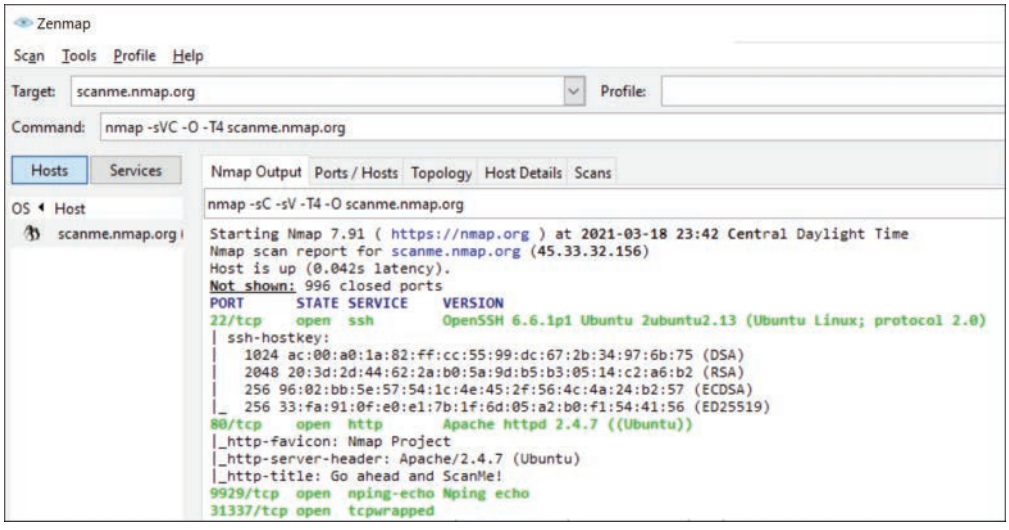


FIGURE 26-4 **nmap** Output on Windows

One thing you might notice when comparing Figure 26-4 and Figure 26-5 is the time to perform the scan and the host latency; in this case, Linux has a slight edge.

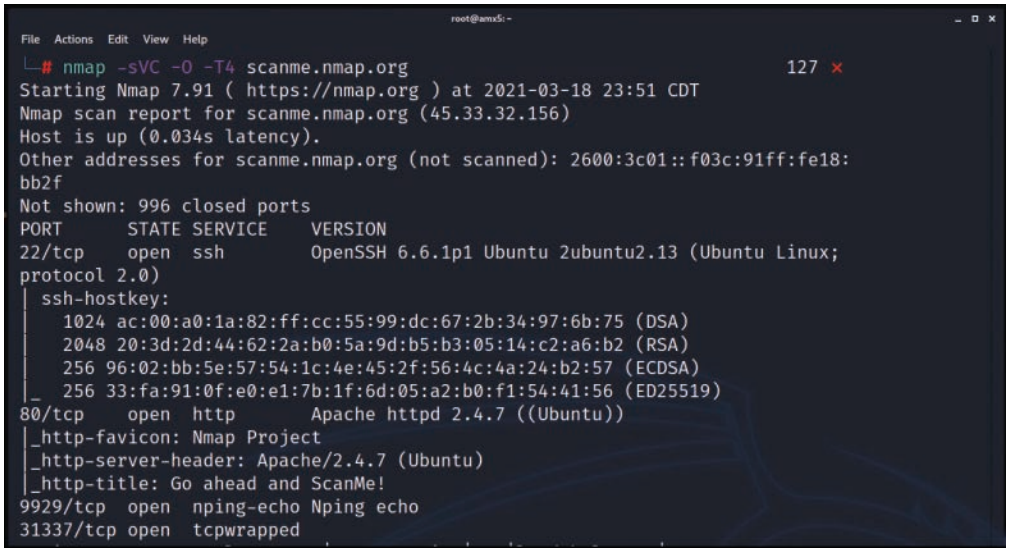


FIGURE 26-5 **nmap** on Linux

The nmap program is a network mapper that has emerged as one of the most diverse, free network discovery tools on the market. It is now one of the core tools

that network administrators use to map their networks. This program can be used to find live hosts on a network and perform port scanning, ping sweeps, OS detection, and version detection.

A number of recent cyber attacks have refocused attention on the type of network auditing that nmap provides. Analysts have pointed out that the recent Capital One hack, for instance, could have been detected sooner if system administrators had been monitoring connected devices. Ideally, nmap should be used as part of an integrated data security program.

ping/pathping

Key Topic

ping (packet internet groper) is the primary TCP/IP command used to troubleshoot connectivity, reachability, and name resolution. Used without parameters, this command displays Help content.

ping verifies IP-level connectivity to another TCP/IP computer by sending Internet Control Message Protocol (ICMP) echo request messages. The corresponding echo reply messages are displayed, along with round-trip times.

You can also use this command to test both the computer name and IP address of the computer. If pinging the IP address is successful, but pinging the computer name isn't, you might have a name resolution problem. In this case, you should make sure the computer name you are specifying can be resolved through the local Hosts file, by using DNS queries, or through NetBIOS name resolution techniques. The command syntax and options/parameters for the **ping** command are as follows:

```
ping [/t] [/a] [/n <count>] [/l <size>] [/f] [/I <TTL>] [/v <TOS>]
[/r <count>] [/s <count>] [{/j <hostlist> | /k <hostlist>}]
[/w <timeout>] [/R] [/S <Srcaddr>] [/4] [/6] <targetname>
```

Example 26-2 demonstrates results from a ping.

Example 26-2 ping Results

```
C:\>ping example.microsoft.com
      Pinging example.microsoft.com [192.168.239.132] with 32 bytes
of data
     Reply from 192.168.239.132:  bytes=32  time=101ms  TTL=124
     Reply from 192.168.239.132:  bytes=32  time=100ms  TTL=124
     Reply from 192.168.239.132:  bytes=32  time=120ms  TTL=124
     Reply from 192.168.239.132:  bytes=32  time=120ms  TTL=124
```

On Linux, the **ping** command works similarly to the way it works in Windows. When you try to ping a remote host, your machine starts sending ICMP echo

requests and waits for a response. If the connection is established, you receive an echo reply for every request.

The output for this command contains the amount of time it takes for every packet to reach its destination and return. The terminal keeps printing the responses until interrupted. After every session, there are a few lines with ping statistics. Figure 26-6 shows the **ping** command on a Linux host, and the list that follows describes what each response represents.

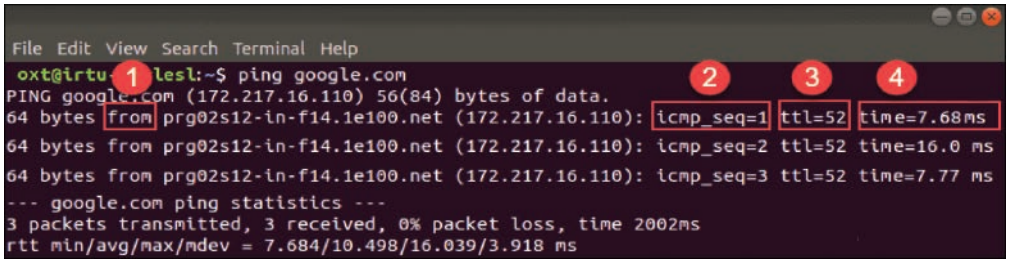


FIGURE 26-6 **ping** Results on a Linux Host

1. **from:** The destination and its IP address.
2. **icmp_seq=1:** The sequence number of each ICMP packet. It increases by one for every subsequent echo request.
3. **ttl=52:** The time-to-live (TTL) value from 1 to 255. It represents the number of network hops a packet can take before a router discards it.
4. **time=7.68 ms:** The time it took a packet to reach the destination and come back to the source, expressed in milliseconds.

Table 26-3 describes the different **ping** options at your disposal.

Table 26-3 **ping** Options

Command Option	Description
a	Generates a sound when the peer can be reached.
b	Allows a ping to broadcast the IP address.
B	Prevents the ping from changing the source address of the probe.
c *count*	Limits the number of sent ping requests.
d	Sets the SO-DEBUG option on the used socket.
f	Floods the network by sending hundreds of packets per second.

Command Option	Description
i *interval*	Specifies an interval between successive packet transmissions. The default value is one second.
I *interface address*	Sets the source IP address to the specified interface address. The option is required when pinging an IPv6 link local address. You can use an IP address or name of the device.
l *preload*	Defines the number of packets to send without waiting for a reply. To specify a value higher than 3, you need superuser permissions.
n	Displays IP addresses in the ping output rather than hostnames.
q	Shows a quiet output. One ping line is displayed and the summary of the **ping** command is output at the end.
T *ttl*	Sets the time to live (TTL).
v	Provides verbose output.
V	Displays the ping version and exits to a new command prompt line.
w *deadline*	Specifies a time limit before the **ping** command exits, regardless of how many packets have been sent or received.
W *timeout*	Determines the time, in seconds, to wait for response.

To ping with **ping6**, you would similarly perform the following command on a Linux host.

You can use the **host** command to find out if a particular host has a corresponding IPv6 address:

```
Linuxhost#-$ host -t AAAA www.google.com
www.google.com has ipv6 address 2607:f8b0:4000:802::1013
```

So then you perform the IPv6 ping using **ping6** as follows:

```
Linuxhost#-$ ping6 -c 4 2607:f8b0:4000:805::1010
64 bytes from 2607:f8b0:4000:805::1010: icmp_seq=1 ttl-51
time=39.5 ms
```

pathping is somewhat of a cross between the ping and tracert utilities. The pathping utility sends packets to each router on the way to a final destination over a period of time and computes results based on the packets that return from each hop. Because pathping shows the degree of packet loss at any given router or link, you can determine which routers or links might be causing network problems. The **pathping** command syntax is as follows:

```
pathping \[-n\] \[-h maximum_hops\] \[-g host-list\] \[-p period\]
\[-q num_queries\] \[-w timeout\] \[-T\] \[-R\] target_name<br>
```

When **pathping** is run, the first results list the path. Next, a busy message is displayed for approximately 90 seconds (the time varies by hop count). During this time, information is gathered from all routers previously listed and from the links between them. At the end of this period, the test results are displayed. For example, the results would be **This Node/Link, Lost/Sent = Pct**, and address columns show that the link between 172.16.87.218 and 192.168.52.1 is dropping 13 percent of the packets. The routers at hops 2 and 4 are also dropping packets addressed to them, but this loss doesn't affect their ability to forward traffic that isn't addressed to them.

The loss rates displayed for the links, identified as a vertical bar (|) in the address column, indicate link congestion that is causing the loss of packets that are being forwarded on the path. The loss rates displayed for routers (identified by their IP addresses) indicate that these routers might be overloaded.

Key Topic

hping

hping is a free TCP/IP packet generator, assembler, and analyzer created by Salvatore Sanfilippo. Although it is similar to the ping utility, it has more functionality than sending a simple ICMP echo request, which ping is usually used for. **hping** can be used to send large volumes of TCP traffic at a target while spoofing the source IP address, making it appear random or even originating from a specific user-defined source.

hping supports TCP, UDP, ICMP and RAW-IP protocols; has a traceroute mode; can send files between a covered channel; and provides many other features. **hping** has a wide range of additional uses, including firewall testing and manual path maximum transmission unit (MTU) discovery. MTU is the size of the largest protocol data unit (packet) that can be communicated in a single network layer transaction. **hping** also has advanced traceroute, remote OS fingerprinting, advanced port scanning, remote uptime guessing, and TCP/IP stack auditing.

You can shore up security when an attacker crafts a packet by using **hping** with the reset flag set and a specified range of ports to scan. This way, the hacker gets no response but can assume the host is alive. If the attacker receives an "icmp host unreachable" message, he or she knows the IP address is not in use.

TIP Hping.org is the home website for the **hping** tool, which contains additional uses and capabilities.

Key Topic

netstat

netstat is a command-line tool in both Windows and Linux. The **netstat** command generates a display that shows network status and protocol statistics. You can display the status of TCP and UDP endpoints in table format, routing table information, and interface information. **netstat** is used on many different operating systems; it is available for Windows, macOS, and Linux.

The proper syntax for the **netstat** command is as follows:

```
netstat [-m] [-n] [-s] [-i | -r] [-f address_family]
```

Table 26-4 describes common **netstat** options.

Table 26-4 netstat Options

Command Option	Description
-a	Displays all connections and listening ports.
-b	Displays the executable involved in creating each connection or listening port. In some cases, well-known executables host multiple independent components, and in these cases, the sequence of components involved in creating the connection or listening port is displayed. In this case, the executable name is in [ ] at the bottom, on top is the component it called, and so forth until TCP/IP was reached. Note that this option can be time consuming and will fail unless you have sufficient permissions.
-e	Displays Ethernet statistics. This option may be combined with the **-s** option.
-n	Displays addresses and port numbers in numerical form.
-o	Displays the owning process ID associated with each connection.
-p proto	Shows connections for the protocol specified by *proto*; *proto* may be any of TPC, UDP, TCPv6, or UDPv6. If used with the **-s** option to display per protocol statistics, *proto* may be any of IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, or UDPv6.
-r	Displays the routing table.
-s	Displays per protocol statistics. By default, statistics are shown for IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, and UDPv6; the **-p** option may be used to specify a subset of the default.
-v	Displays the sequence of components involved in creating the connection or listening port for all executables when used in conjunction with **-b**.
interval	Redisplays selected statistics, pausing for an interval specified in seconds between each display. Press Ctrl+C to stop redisplaying statistics. If omitted, **netstat** prints the current configuration information once.

The **netstat -s** option displays per protocol statistics for the UDP, TCP, ICMP, and IP protocols. Figure 26-7 shows the results of a **netstat -s** command, excluding some output for brevity. Using the **-s** option can show you where a particular protocol might be having issues.

```
C:\Users\root>netstat -s

IPv4 Statistics

  Packets Received                   = 85521942
  Received Header Errors             = 0
  Received Address Errors            = 3
  Datagrams Forwarded                = 0
  Unknown Protocols Received         = 0
  Received Packets Discarded         = 84340
  Received Packets Delivered         = 85642843
  Output Requests                    = 29465709
  Routing Discards                   = 0
  Discarded Output Packets           = 9
  Output Packet No Route             = 116
  Reassembly Required                = 0
  Reassembly Successful              = 0
  Reassembly Failures                = 0
  Datagrams Successfully Fragmented  = 0
  Datagrams Failing Fragmentation    = 0
  Fragments Created                  = 0
```

FIGURE 26-7 Using **netstat -s** to Determine Where Protocols Might Be Having Issues

Figure 26-8 demonstrates another really useful and common **netstat** command—the **netstat -an** command, which displays all connections and listening ports and displays addresses and port numbers in numerical format.

```
Command Prompt
C:\Users\root>netstat -an

Active Connections

  Proto  Local Address          Foreign Address        State
  TCP    0.0.0.0:135            0.0.0.0:0              LISTENING
  TCP    0.0.0.0:443            0.0.0.0:0              LISTENING
  TCP    0.0.0.0:445            0.0.0.0:0              LISTENING
  TCP    0.0.0.0:808            0.0.0.0:0              LISTENING
  TCP    0.0.0.0:903            0.0.0.0:0              LISTENING
  TCP    0.0.0.0:913            0.0.0.0:0              LISTENING
  TCP    1.1.100.45:51936       72.21.91.29:80         CLOSE_WAIT
  TCP    1.1.100.45:52750       3.235.71.133:443       CLOSE_WAIT
  TCP    1.1.100.45:52862       52.96.16.162:443       ESTABLISHED
  TCP    1.1.100.45:52863       52.96.57.50:443        ESTABLISHED
  TCP    1.1.100.45:52865       52.96.57.50:443        ESTABLISHED
```

FIGURE 26-8 **netstat -an** Displays Host Connections

NOTE The **netstat** command-line utility displays all ports a computer is listening on and is useful for quickly determining active connections. It can also display the routing table and preprotocol statistics. Open ports on a system provide entryways for attackers to exploit and gain control. You also can use **netstat** as a forensics tool to see what programs or processes are active on a system.

Key Topic

netcat

The ***netcat*** utility program supports a wide range of commands to manage networks and monitor the flow of traffic data between systems. You can think of it as a free and easy companion tool to use alongside Wireshark (discussed shortly), which specializes in the analysis of network packets. The original version of **netcat** was released in 1995 and has received a number of iterative updates in the decades since. The **netcat** tool is available in Windows and Linux, although the Linux version works more reliably.

netcat can be a useful tool for any security team, although the growth of internally managed network services and cloud computing make that particular environment a good fit. Network and system administrators need to be able to quickly identify how their network is performing and what type of activity is occurring. **netcat** functions as a back-end tool that allows for port scanning and port listening. In addition, you can actually transfer files directly through **netcat** or use it as a backdoor into other networked systems.

After you have the netcat application set up on your Windows or Linux server, you can start running basic commands to test its functionality and get a better understanding of how an attacker would use it to infiltrate your network. Here are a few to get started with:

- **nc -help:** This command prints a list of all of the available commands.
- **nc -z -v abc.com:** This command runs a basic port scan of the specified abc.com site or server. **netcat** returns verbose results with lists of ports and statuses. Keep in mind that you can use an IP address in place of the site domain.
- **nc -l:** This command instructs the local system to begin listening for TCP connections and UDP activity on a specific port number.
- **nc site.com 1234 (less than)** *file_name***:** This command initiates the transfer of a file based on the specified port number.
- **Printf: netcat** can actually operate as a simplified web host. This command lets you save HTML code and publish it through your local server.

When you're trying to diagnose a network issue or performance problem, executing a port scan with netcat is a smart first step to take. The scan checks the status of all ports on the given domain or IP address so that you can determine whether a firewall or other blocking mechanism is in place.

A basic port scan command for an IP address looks like this:

```
nc -v -n 8.8.8.8 1-1000
```

Key Topic

IP Scanners

IP scanners are tools that are quite useful in the field of networking. An IP scanner is, as its name indicates, a device that scans for IP addresses and various other information about the devices on your network. In short, the IP scanner scans your network for devices and information relevant to them.

The first—and a very important—reason to use an IP scanner is for security purposes. You can check the devices that are connected to your network. You also get detailed information about the devices on the network you scan. This information can help you keep an eye on the devices and help you look out for unknown or suspicious devices on the network. Different IP scanners provide different levels of information about specific hosts. Typically, they include IP address, MAC address, vendor, operating system, number of open TCP/UDP ports, status of ports, and a description of the host.

Modern IP scanners can also determine the patch level of a host, users who are logged in, and much more. Automated patching systems use a type of IP scanner to help provide this type of information.

A number of IP scanners are freely available, including the SolarWinds IP scanner, Angry IP scanner, Advanced IP scanner, and PRTG network monitor.

Key Topic

arp

The **arp** command is a TCP/IP utility and Microsoft Windows command for viewing and modifying the local Address Resolution Protocol (ARP) cache, which contains recently resolved MAC addresses of Internet Protocol (IP) hosts on the network. When one host on a TCP/IP network wants to communicate with another host, the first host begins by using ARP to resolve the IP address of the other host into its associated MAC address. The MAC address is needed for communication to take place over the network. The **arp ping** command uses ARP to test connectivity between systems rather than using ICMP, as done with a regular **ping**.

Figure 26-9 shows the **arp** options on a Windows host. The most common is **arp -a**, which displays all ARP entries recently discovered.

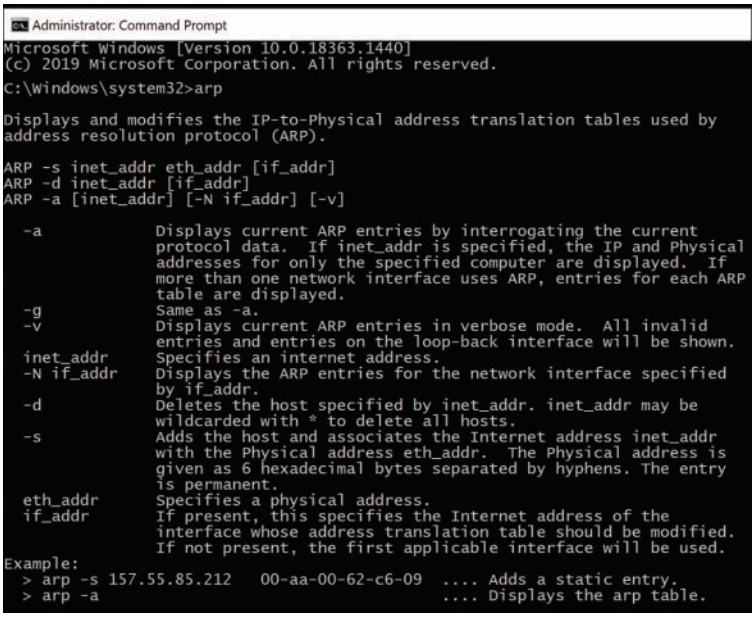


FIGURE 26-9 **arp** Command Options

Table 26-5 shows the command options available for **arp**.

Table 26-5 **arp** Options

Command Option	Description
-a **-g**	Displays current ARP entries by interrogating the current protocol data
-v	Displays current ARP entries in verbose mode; also displays all invalid entries and loopback interfaces
inet_addr	Specifies an Internet address
-N if_addr	Displays the ARP entries for the network interface specified by the **if_addr**
-d	Deletes the host specified by **inet_addr**
-s	Adds the host and associates the Internet address **inet_addr** with the physical address

For example, to add a static entry, you would enter the following command:

```
arp -s 192.168.0.0 00-aa-00-62-c6-09
```

NOTE Address Resolution Protocol (ARP) poisoning is limited to attacks that are locally based, so an intruder needs either physical access to your network or control of a device on your local network to perform such attacks.

Key Topic

route

route is a command used to view and manipulate the IP routing table in Microsoft Windows operating systems and also in Linux and other Linux-type systems like routers and switches. Manual manipulation of the routing table is characteristic of static routing. The **route** command operates differently under Windows than it does in Linux, so learning the differences is important.

In Linux distributions based on 2.2.x Linux kernels, the **route** command operated to connect a computer to a network and to define routes between computer networks. Distributions based on later kernels have deprecated **route**, replacing it with **iproute2**. For a typical computer that has a single network interface and is connected to a local area network that has a router, the routing table is pretty simple and isn't often the source of network problems.

If you're having trouble accessing other computers or other networks, you can use the **route** command to make sure that a bad entry in the computer's routing table isn't the culprit. For a computer that has more than one interface and that is configured to work as a router, or the host needs access to two separate networks, the routing table is often a major source of trouble. Setting up the routing table properly is a key part of configuring a router to work.

To display the routing table (both IPv4 and IPv6) in Windows, you can use the **route print** command. In Linux, you can just use **route** without any command-line switches. The output displayed by the Windows and Linux commands is similar. Here's an example from a typical Windows client computer:

```
C:\Users\jenny>route print
===========================================================================
Interface List
22...00 4f 02 b7 34 ea ......Killer E2400 Gigabit Ethernet Controller
12...42 5b dd ea af fb ......Microsoft Wi-Fi Direct Virtual Adapter
===========================================================================
IPv4 Route Table
===========================================================================
Active Routes:
Network Destination         Netmask          Gateway       Interface   Metric
          0.0.0.0           0.0.0.0      12.11.100.1       12.11.100.1      25
      12.11.100.0     255.255.255.0          On-link       12.11.100.1     281
    12.11.100.255   255.255.255.255          On-link       12.11.100.1     281
```

For each entry in the routing table, five items of information are listed:

- Destination IP address, which is actually the address of the destination subnet and must be interpreted in the context of the subnet mask
- The subnet mask that must be applied to the destination address to determine the destination subnet

- The IP address of the gateway to which traffic intended for the destination subnet will be sent
- The IP address of the interface through which the traffic will be sent to the destination subnet
- The metric, which indicates the number of hops required to reach destinations via the gateway

Each packet that's processed by the computer is evaluated against the rules in the routing table. If the packet's destination address matches the destination subnet for the rule, the packet is sent to the specified gateway via the specified network interface.

TIP One major difference between the Windows version and Linux versions of **route** is the order in which they list the routing table. The Windows **route** command lists the table starting with the most general entry and works toward the most specific. The Linux version is the other way around.

You can make sure your systems stay secure by limiting who has access to execute the **route** command on them.

Key Topic

curl

curl (or cURL) is a command-line tool used to transfer data to or from a server, using any of the supported protocols: HTTP, FTP, IMAP, POP3, SCP, SFTP, SMTP, TFTP, TELNET, LDAP, or FILE. It is powered by Libcurl. This tool is preferred for automation because it is designed to work without user interaction. It can transfer multiple files at once.

The syntax for the **curl** command is as follows:

```
curl [options] [URL...]
```

This command supports HTTPS and performs SSL certificate verification by default when a secure protocol is specified such as HTTPS. When it connects to a remote server via HTTPS, it obtains the remote server certificate and then checks the validity of the certificate and remote server against its CA certificate store to ensure the remote server is the one it claims to be. Some curl packages are bundled with the CA certificate store file.

When you try to access an SSL/TLS cert-secured URL, if it has the wrong cert or the CN doesn't match, you get the following error:

```
curl: (51) Unable to communicate securely with peer: requested domain
name does not match the server's certificate.
```

To ignore the certificate, you can use **curl --insecure https://yoursite.com** as an option to skip certificate verification. To use SSL/TLS and connect using SSL v3, you would enter **curl -sslv3 https://yourside.com**.

NOTE The **curl** tool can get and send data using URLs. For example, typing the command followed by any URL returns the HTML source for the web page. This tool can also be used to grab banner information from websites, including cookies, headers, and user-agent information.

Key Topic

theHarvester

theHarvester is a simple but effective tool in every pen tester's toolchest. It was developed in Python. TheHarvester can gather information like emails, subdomains, hosts, employee names, open ports, and banners from different public sources like search engines, PGP key servers, LinkedIn, Twitter, and there's a plug-in for the Shodan database. Shodan is a searchable database of Internet-connected devices that were discovered and stored by Shodan.

Ideally, you want to run this tool against your organization and determine which users/accounts are made public. You should also encourage employees to not use their work email address for public posting.

Figure 26-10 shows theHarvester interface. When you enter **theHarvester** at the command line, it provides a list of options. The main syntax is as follows:

```
theHarvester [-h] -d DOMAIN [-l LIMIT] [-S START] [-g][-g][-s]
```



FIGURE 26-10 theHarvester Command-line Interface

One option is **-d** for domain:

```
theHarvester -d mydomain.org -l 500 -b google
```

Key Topic

sn1per

The ***sn1per*** tool is an automated scanner that you can use during penetration testing to perform vulnerability scanning. Two sn1per versions are available: the Community Edition and the Professional Edition.

The current list of capabilities includes automatically performing basic reconnaissance (i.e., whois, ping, DNS), launching Google hacking queries against a target domain, enumerating open ports, brute-forcing subdomains and DNS information, checking for subdomain hijacking, running targeted nmap scripts against open ports, running targeted Metasploit scan and exploit modules, scanning all web applications for common vulnerabilities, brute-forcing all open services, exploiting remote hosts to gain remote shell access, and performing high-level enumeration of multiple hosts. Another feature called Autopwn was added for Metasploitable, ShellShock, MS08-067, and Default Tomcat Creds.

Figure 26-11 illustrates the sn1per command line and options; you need to specify a target or enter **-help** for command usage. About 100 options are available; the main one is **sniper -t** <*TARGET*>, or you can use a custom config file for your environment by selecting **sniper -C /full/path/to/sniper.conf -t** <*TARGET*> **-m** <*MODE*> **-w** <*WORKSPACE*>.

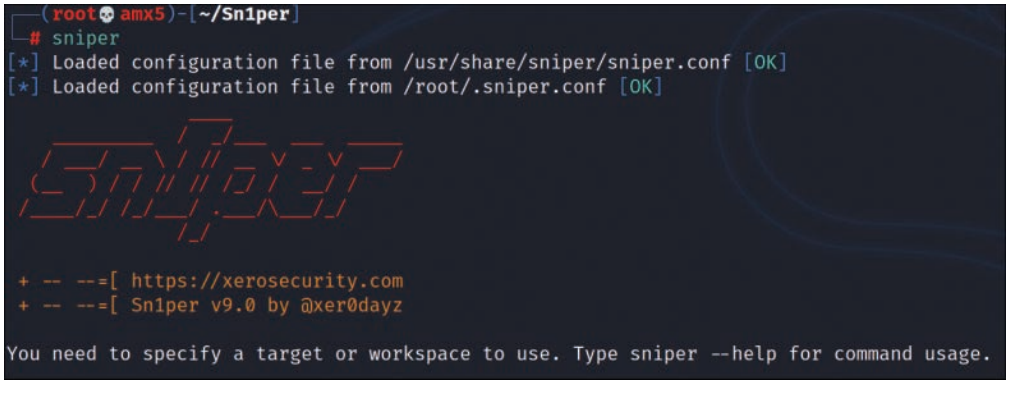


FIGURE 26-11 sn1per Command Line and Options

There is also a stealth osint recon mode: **sniper -t** <*TARGET*> **-m stealth -o -re**.

To obtain sn1per help, you enter **sniper -help.**

Figure 26-12 shows sn1per launched against evilcorp.biz in stealth mode.

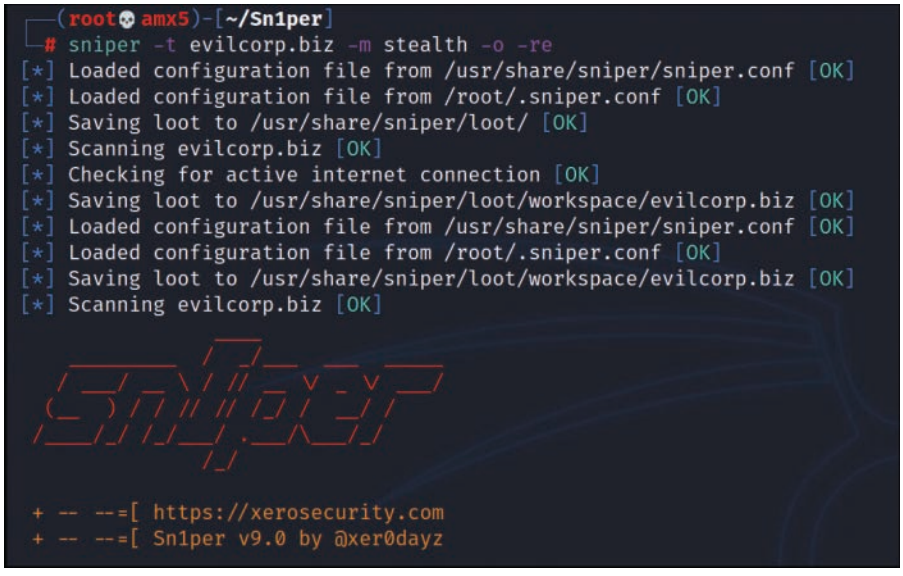


FIGURE 26-12 sn1per in Stealth Mode

Key Topic

scanless

scanless is a Python 3 Linux command-line utility for exploitation websites. This utility functions as a public port scan scraper; it can use a number of public scan sites to perform port scans on your behalf. It is useful for early stages of penetration tests when you'd like to run a live port scan on a host without having it originate from your IP address. There also are other public port scanners that you can use to scan devices so that they don't appear to be coming from your IP address; these include YouGetSignal, Ping.eu, ViewDNS, IPFingerPrints, and many others.

Figure 26-13 shows the command-line switches and options for scanless. A common option is the **-t** command for *target*, as demonstrated here:

```
scanless -t scanme.nmap.org -s ipfingerprints
```

```
┌──(root💀amx5)-[~/Sn1per]
└─# scanless
usage: scanless [-h] [-v] [-t TARGET] [-s SCANNER] [-r] [-l] [-a] [-d]

scanless, an online port scan scraper.

optional arguments:
  -h, --help            show this help message and exit
  -v, --version         display the current version
  -t TARGET, --target TARGET
                        ip or domain to scan
  -s SCANNER, --scanner SCANNER
                        scanner to use (default: hackertarget)
  -r, --random          use a random scanner
  -l, --list            list scanners
  -a, --all             use all the scanners
  -d, --debug           debug mode (cli mode off & show network errors)
```

FIGURE 26-13 scanless Command Switches and Options

Key Topic

dnsenum

dnsenum is a command-line tool that automatically identifies basic DNS records such as MX (mail exchange servers), NS (domain name servers), or A (the address record for a domain). It also attempts zone transfers on all identified servers and can attempt reverse resolution (that is, getting the hostname given an IP address) and brute-forcing (querying for the existence of hostnames to get their IP address) subdomains and hostnames.

The syntax for the **dnsenum** command is as follows:

```
dnsenum -enum <url>
```

The tool starts querying the DNS servers of domain.org. The first information that you get is the host address, which for scanme.org is x.x.x.x. Next, you see the name servers, which give you an idea of the hosting provider that domain.org is using. After that, you get the MX record, where you can see the mail server of your target host.

After that, dnsenum starts the DNS zone transfer. When performing a zone transfer, you can discover more information about a domain, such as any subdomains that are included in the same zone and SOA records. Figure 26-14 shows the results after performing a zone transfer for a domain that does not have a DNS server listening.

```
┌──(root💀amx5)-[~/Sn1per]
└─# dnsenum scanme.nmap.org
dnsenum VERSION:1.2.6

-----   scanme.nmap.org   -----

Host's addresses:
__________________

scanme.nmap.org.                         3600     IN    A        45.33.32.156

Name Servers:
_____________

 scanme.nmap.org NS record query failed: NOERROR
```

FIGURE 26-14 An Invalid Transfer for the Domain Using **dnsenum**

Figure 26-15 shows the results after performing a valid zone transfer.

```
┌──(root💀amx5)-[~/Sn1per]
└─# dnsenum zonetransfer.me
dnsenum VERSION:1.2.6

-----   zonetransfer.me   -----

Host's addresses:
__________________

zonetransfer.me.                         7200     IN    A        5.196.105.14

Name Servers:
_____________

nsztm1.digi.ninja.                       10800    IN    A        81.4.108.41
nsztm2.digi.ninja.                       10800    IN    A        34.225.33.2

Mail (MX) Servers:
__________________

ASPMX3.GOOGLEMAIL.COM.                   215      IN    A        173.194.209.26
ASPMX2.GOOGLEMAIL.COM.                   293      IN    A        142.250.125.26
ALT1.ASPMX.L.GOOGLE.COM.                 269      IN    A        142.250.125.26
ASPMX5.GOOGLEMAIL.COM.                   293      IN    A        64.233.177.26
ASPMX.L.GOOGLE.COM.                      199      IN    A        142.250.138.26
ASPMX4.GOOGLEMAIL.COM.                   293      IN    A        142.250.96.26
ALT2.ASPMX.L.GOOGLE.COM.                 12       IN    A        173.194.209.27
```

FIGURE 26-15 A Valid Zone Transfer from a Domain



Nessus

Nessus was built from the ground up with a deep understanding of how security practitioners work. This remote security scanning tool scans a computer and notifies you if it discovers any vulnerabilities that malicious hackers could use to gain access to any computer you have connected to a network. It does this by running over 1200 checks on a given computer, testing to see if any of these attacks could be used to break into the computer or otherwise harm it.

Most high-level network traffic, such as email and web pages, reaches a server via a high-level protocol that is transmitted reliably by a TCP stream or unreliably via UDP. To keep different streams from interfering with each other, a computer divides its physical connection to the network into thousands of logical paths, called *ports*. So, if you want to talk to a web server on a given machine, you would connect to port 443 (the standard HTTPS secure port), but if you want to connect to an SMTP server on that same machine, you would instead connect to port 25.

Each computer has thousands of ports, all of which may or may not have services running or listening on them. Nessus works by testing each port on the computer, determining what service it is running, and then testing this service to make sure it has no vulnerabilities that could be used by hackers to carry out a malicious attack, and provides a report of ports that may require attention.

Nessus is called a "remote scanner" because it does not need to be installed on the computer you plan to test. Instead, you should install it on only one computer and then can test as many computers as you would like. Nessus was developed on the Linux platform but is now available on Windows as well. The open-source version of Nessus is called Open Source Vulnerability Assessment (OpenVAS). On a Linux host where you would like to start Nessus, you can simply enter **nessusd start** to start the Nessus daemon. Then you start your web browser and connect to the server.

Figure 26-16 shows a Nessus scan being conducted via the web browser connected to a Linux host running Nessus. This example uses an essentials license that allows you to scan 12 hosts.

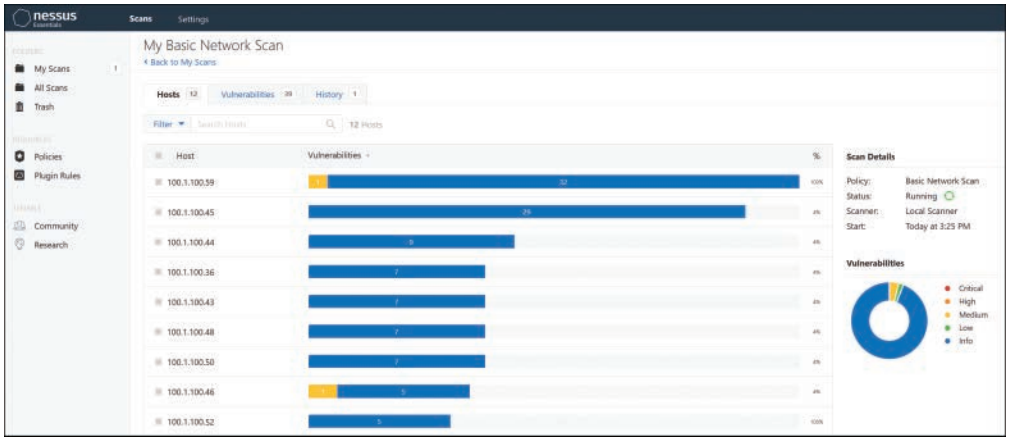


FIGURE 26-16 Nessus Interface via Web Browser

The steps you follow during scanning can be summarized as follows:

Step 1. Define your scan parameters.

Step 2. Create a scan.

Step 3. Launch the scan.

Step 4. Analyze the scan results.

Key Topic

Cuckoo

Cuckoo Sandbox is free, open-source software for automating analysis of suspicious files. To do so, it makes use of custom components that monitor the behavior of the malicious processes while running in an isolated environment. You can throw any suspicious file at it, and in a matter of minutes, Cuckoo will provide a detailed report outlining the behavior of the file when executed inside a realistic but isolated environment.

Malware is the Swiss army knife of cybercriminals and any other adversary to your corporation or organization. In these evolving times, detecting and removing malware artifacts is not enough: it's vitally important to understand how they operate in order to understand the context, the motivations, and the goals of a breach.

Cuckoo Sandbox automates the task of analyzing any malicious files under Windows, macOS, Linux, and Android.

Because Cuckoo is open source and has an extensive modular design, you can customize any aspect of the analysis environment, analysis results processing, and reporting stage. Cuckoo provides all the requirements to easily integrate the

sandbox into your existing framework and back end in the way you want, with the format you want, and all of that without licensing requirements.

Figure 26-17 shows a Cuckoo 2.0.0 analysis results page with the side bar locked in permanently open status.

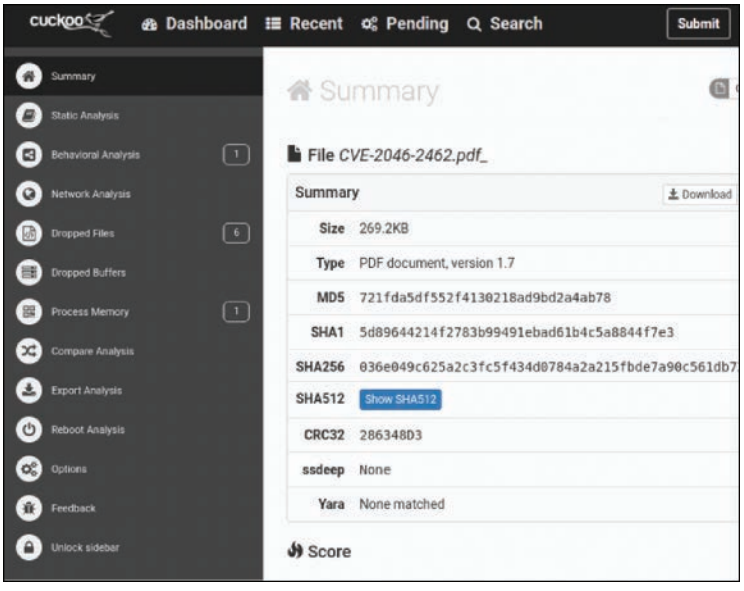


FIGURE 26-17 Cuckoo GUI

Cuckoo can retrieve the following types of results:

- Traces of calls performed by all processes spawned by malware
- Files being created, deleted, and downloaded by malware during its execution
- Memory dumps of malware processes
- Network traffic traces in PCAP format
- Screenshots taken during the execution of malware
- Full memory dumps of the machines

File Manipulation

Manipulating files without it being noticed is a trick of the hackers' trade. Most applications/programs require the ability to read or write a file. The operating system gives permission to the program for an operation on the file and provides an

interface to the user to create or delete files and directories. The operating system provides an interface to create a backup of the file system.

A file represents a collection of related information. Computers may store files on a disk media (secondary storage) for long-term storage purposes. Examples include magnetic tape, magnetic disk, and optical disk drives like CDs and DVDs. They have their own properties like speed, capacity, data transfer rate, and data access methods.

A file system is generally organized into directories for easy navigation and usage. Some of these directories can contain files and other directories. The major activities of an operating system with respect to file management should be to allow a program to read or write a file. These permissions include read-only, read-write, denied, and so on.

Key Topic

head

The ***head*** command is a Linux-centric command that reads the first 10 lines of any given filename. The basic syntax of the **head** command is as follows:

```
head [options] [file(s)]
```

For example, the following command displays the first 10 lines of the file named /etc/passwd:

```
head /etc/passwd
```

If more than one file is given, **head** shows the first 10 lines of each file separately.

If you want to retrieve more than the default 10 lines, you can use the **-n** option along with an integer telling the number of lines to be retrieved. Because this tool can display sensitive information, it is best to limit access to authorized administrators only. Figure 26-18 shows the **head /etc/passwd** output on a Linux host. As you can see, there is the root account. If you want to view only two lines, you would enter **head /etc/passwd -n 2**.

```
└─# head /etc/passwd
root:x:0:0:root:/root:/usr/bin/zsh
daemon:x:1:1:daemon:/usr/sbin:/usr/sbin/nologin
bin:x:2:2:bin:/bin:/usr/sbin/nologin
sys:x:3:3:sys:/dev:/usr/sbin/nologin
sync:x:4:65534:sync:/bin:/bin/sync
games:x:5:60:games:/usr/games:/usr/sbin/nologin
man:x:6:12:man:/var/cache/man:/usr/sbin/nologin
lp:x:7:7:lp:/var/spool/lpd:/usr/sbin/nologin
mail:x:8:8:mail:/var/mail:/usr/sbin/nologin
news:x:9:9:news:/var/spool/news:/usr/sbin/nologin
```

FIGURE 26-18 **head /etc/passwd** Output

tail

The ***tail*** command allows you to display the last 10 lines of any text file. Similar to the **head** command, the **tail** command also supports options for *n* number of lines and *n* number of characters. It is one of the most commonly used tools in troubleshooting.

The basic syntax of the **tail** command is as follows:

```
tail [options] [filenames]
```

For example, the following command prints the last 10 lines of a file called access.log:

```
tail access.log
```

If more than one file is provided, **tail** prints the last 10 lines of each file. It is a great tool to continuously monitor log files, which are always changing. You simply enter **tail -f** *file* as shown in Figure 26-19.

```
└─# tail -f /var/log/auth.log
Mar 19 16:17:01 amx5 CRON[3222]: pam_unix(cron:session): session closed for
er root
Mar 19 16:20:51 amx5 lightdm: pam_unix(lightdm-greeter:session): session ope
d for user lightdm by (uid=0)
Mar 19 16:20:51 amx5 systemd-logind[529]: New session c6 of user lightdm.
Mar 19 16:20:51 amx5 systemd: pam_unix(systemd-user:session): session opened
or user lightdm by (uid=0)
```

FIGURE 26-19 **tail -f** Output

Administrators often use the **tail** command to monitor log files associated with security components. It only views and does not modify files.

cat

The Linux ***cat*** command is a widely used, universal tool. It copies standard input to standard output. The command supports scrolling if the text file doesn't fit the current screen.

The basic syntax of the **cat** command is as follows:

```
cat [options] [filenames] [-] [filenames]
```

The most frequent use of this command is to read the contents of files. All that is required to open a file for reading is to type **cat** followed by a space and the filename. For example, if you want to view the /etc/passwd file, you would enter

```
$ cat /etc/passwd
```

This command can be also used to create files. This result is achieved by executing **cat** followed by the output redirection operator and the filename to be created. To make a file called newfile.txt, you would enter **$ cat > newfile.txt**.

grep

Global Regular Expression Print (***grep***) is a Linux program for finding matching patterns, to search for a string of characters in a specified file. It's like the Swiss army knife of file tools. It can be found on Linux as well as macOS. It can read just about any text, meaning it can read input from other commands, or it can open and look through files directly. This tool is incredibly useful, especially for looking through directories from the command line.

The **grep** command consists of three parts in its most basic form. The first part starts with **grep**, followed by the pattern that you're searching for. After the string comes the filename that the tool searches through. The simplest **grep** command syntax looks like this:

```
$ grep host sample
```

This command displays every line where there is a match for the word *host*. When executing this command, you do not get exact matches. Instead, the terminal prints the lines with words containing the string of characters you entered. The command can contain many options, pattern variations, and filenames. You can combine as many options as necessary to get the results you need.

To search multiple files with the **grep** command, you insert the filenames you want to search, separated with a space character. In this case, you can use the **grep** command to match a word:

```
$ grep host sample sample2 sample3
```

The terminal prints the name of every file that contains the matching lines, and the actual lines that include the required string of characters. Because **grep** commands are case sensitive, one of the most useful operators for **grep** searches is **-i**. Instead of printing lowercase results only, the terminal displays both uppercase and lowercase results. The output includes lines with mixed case entries.

grep can also use regular expressions such as this: **grep "\<[A-Za-z].*\>"** *filename*.

Table 26-6 describes the Linux **grep** command options.

Key Topic

Table 26-6 grep Options

Command Option	Description
-i	Ignores case distinctions on Linux
-w	Forces a pattern to match only whole words
-v	Selects nonmatching lines
-n	Prints the line number with output lines
-h	Suppresses the Linux filename prefix on output
-r	Searches directories recursively on Linux
-R	Searches just like **-r** but follows all symlinks (symbolic links)
-l	Prints only names of files with selected lines
-c	Prints only a count of selected lines per file
--color	Displays matched pattern in colors

Key Topic

chmod

In Linux operating systems, ***chmod*** is the command and system call that is used to change the access permissions of file system objects. It is also used to change special mode flags. The name is an abbreviation of *change mode*. The request is filtered by the umask capability. **Umask** is a command that determines the settings of a mask that controls how file permissions are set for newly created files.

In Linux, who can do what to a file or directory is controlled through sets of permissions. There are three sets of permissions: one set for the owner of the file, another set for the members of the file's group, and a final set for everyone else. The permissions control the actions that can be performed on the file or directory. They either permit, or prevent, a file from being read, modified, or if it is a script or program, executed. For a directory, the permissions govern who can change directory (**cd**) and who can create or modify files within the directory. You use the **chmod** command to set each of these permissions.

There are three characters in each set of permissions. These characters are indicators for the presence or absence of one of the permissions. They are either a dash (-) or a letter. The dash character means that permission is not granted. If the character is r, w, or x, that permission has been granted. The letters represent the following:

- **r:** Read permissions. The file can be opened and its content viewed.
- **w:** Write permissions. The file can be edited, modified, and deleted.
- **x:** Execute permissions. If the file is a script or program, it can be run (executed).

For example:

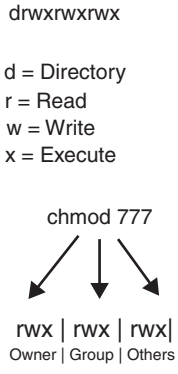
The **---** means no permissions have been granted at all.

The **rwx** means full permissions have been granted. The read, write, and execute indicators are all present.

To use the **chmod** command to set permissions, you need to tell it three things: first, who (who are you setting permissions for?); second, what (what change are you making? Are you adding or removing the permissions?); and third, which (which of the permissions are you setting?)

For example, if you want the user Chris to have ownership of the file as well as read and write permissions, you can use **chown** to change ownership: **chown chris:chris new_file.txt**. Then you can use **chmod u=rw,og=r new_file.txt**.

Figure 26-20 illustrates using **chmod** to change owner, group, and others, as well as what each number correlates to. For example, 777 provides full access to a particular file or directory.



7	rwx	111
6	rw-	110
5	r-x	101
4	r--	100
3	-wx	011
2	-w-	010
1	--x	001
0	---	000

FIGURE 26-20 **chmod** Capabilities

Logger

The ***logger*** command is used to add logs to the local syslog file or a remote syslog server, as follows:

```
logger [-isd] [-f file] [-p pri] [-t tag] [-u socket] [message ...]

logger -n 192.168.1.100
```

Logger can make entries in the system log. "A short log line with some details or context" adds the message in quotation marks to the log of the remote server at the given IP address.

The logger utility exits 0 on success and > 0 if an error occurs. Valid facility names are auth, authpriv (for security information of a sensitive nature), cron, daemon, ftp, kern, lpr, mail, news, security (deprecated synonym for auth), syslog, user, uucp, and local0 to local7, inclusive.

Valid level names are alert, crit, debug, emerg, err, error (deprecated synonym for err), info, notice, panic (deprecated synonym for emerg), warning, and warn (deprecated synonym for warning).

The syntax for the **logger** command is as follows:

```
logger [-isd] [-f file] [-p pri] [-t tag] [-u socket] [message ...]
```

Table 26-7 describes these options and parameters in more detail.

Key Topic

Table 26-7 logger Options

Command Option	Description
-i	Logs the process ID of the logger process with each line.
-s	Logs the message to standard error, as well as the system log.
-f *file*	Logs the specified file.
-p *pri*	Enters the message with the specified priority. The priority may be specified numerically or as a facility.level pair. For example, **-p local3.info** logs the message(s) as *informational level* in the *local3 facility*. The default is "user.notice."
-t *tag*	Marks every line in the log with the specified *tag*.
-u *sock*	Writes to a *sock* as specified with the socket instead of built-in syslog routines.
-d	Uses a datagram instead of a stream connection to this socket.
--	Ends the argument list. This way, the *message* can start with a hyphen (-).
message	Writes the message to the log; if not specified, and the **-f** flag is not provided, standard input is logged.

Key Topic

Shell and Script Environments

Environmental variables are defined for the current shell and are inherited by any child shells or processes. Environmental variables are used to pass information into processes that are spawned from the shell. ***Shell variables*** are contained exclusively within the shell in which they were set or defined. They are often used to keep track of ephemeral data, like the current working directory.

A shell maintains an environment that includes a set of variables defined by the login program, system initialization file, and user initialization files. In addition, some variables are defined by default.

A shell can have two types of variables:

- **Environment variables:** Variables that are exported to all processes spawned by the shell
- **Shell (local) variables:** Variables that affect only the current shell

There are several types of shells, including C, Bourne, and Korn shells. In the Bourne and Korn shells, you can use the uppercase variable name equal to some value to set both shell and environment variables. For all shells, you generally refer to shell and environment variables by their uppercase names. In a user initialization file, you can customize a user's shell environment by changing the values of the predefined variables or by specifying additional variables.

Key Topic

SSH

The Secure Shell (***SSH***) is a cryptographic network protocol. It gives users a secure way to access a computer over an unsecured network that can establish a secure session between a client and host computer using an authenticated and encrypted connection over port 22. SSH encrypts all data, including the login portion.

When you run a shell script on a remote machine over SSH, the script might fail because it depends on environment variables, which are unset on the remote machine. The appropriate command-line syntax is **ssh user@ip "sh abc.sh"**. When you are interacting with your server through a shell session, your shell compiles environment variables and other pieces of information to determine its behavior and access to resources. Some of these settings are contained within configuration or environment settings, and others are determined by user input.

The shell keeps track of all of these settings and details through an area it maintains called the environment. The shell builds the environment every time it starts a session that contains variables that define system properties.

Every time a shell session spawns, a process takes place to gather and compile information that should be available to the shell process and its child processes. It obtains the data for these settings from a variety of different files and settings on the system. The environment provides a medium through which the shell process can get or set settings and, in turn, pass them on to its child processes.

The environment is implemented as strings that represent key-value pairs. If multiple values are passed, they are typically separated by colon (:) characters. Each shell session keeps track of its own shell and environmental variables. You can access them in a few different ways. You can see a list of all of your environmental variables by using the **env** or **printenv** commands.

Some environmental and shell variables are very useful and are referenced fairly often. You likely will come across the following common environmental variables:

- **SHELL:** The shell that interprets any commands you type in. In most cases, this is bash by default, but you can set other values if you prefer other options.
- **TERM:** The type of terminal to emulate when running the shell. Different hardware terminals can be emulated for different operating requirements. You usually don't need to worry about this, though.
- **USER:** The currently logged-in user.
- **PWD:** The current working directory.
- **OLDPWD:** The previous working directory. The shell keeps this directory so that you can switch back to your previous directory by running **cd -**.
- **LS_COLORS:** The color codes that are used to optionally add colored output to the **ls** command. These colors are used to distinguish different file types and provide more information to the user at a glance.
- **MAIL:** The path to the current user's mailbox.
- **PATH:** A list of directories that the system checks when looking for commands. When a user types in a command, the system checks directories in this order for the executable.
- **LANG:** The current language and localization settings, including character encoding.
- **HOME:** The current user's home directory.
- **_:** The most recently previously executed command.



PowerShell

In ***PowerShell***, environment variables are stored in the Env "drive," which is accessible through the PowerShell environment provider, a subsystem of PowerShell. This isn't a physical drive; instead, it's a virtual file system. Environment variables convey information about your login session to your computer. PowerShell has a feature called *providers* that creates one or more drives, which are hierarchical, file-system-like structures that allow a user to manage various areas in Windows. One of those providers is for environment variables called *Environment*. You can also access environment variables using the built-in variable called $env followed by a colon and the name of the environment variable. For example, instead of using Get-Item or Get-ChildItem and using the Env drive, you can save some keystrokes and instead simply specify **$env:COMPUTERNAME**.

Key Topic

Python

The clear, simple syntax of ***Python*** makes it an ideal language to interact with REST APIs, and in typical Python fashion, there's a library made specifically to provide that functionality. Python Requests is a powerful tool that provides the simple elegance of Python to make HTTP requests to any API in the world.

Environment variables help you move away from manually updating your path or other environment settings each time they change. For example, instead of manually changing the "user" part of a path when you run a script on someone else's machine, you can use an environment variable that returns the name of that user. This means you no longer have to remember to change that when using a script on a different system.

One major advantage to using environment variables is when there are aspects of your code you don't want others to see, such as API tokens that grant access to your accounts. If these tokens are stored as environment variables, you don't need to explicitly type them out in the file; you only need to call the environment variables. This also means that other people can use your code without having to hard-code their own API tokens.

Environment variables are implemented through the OS package, specifically os.environ. To see all environment variables on your system, just call it by using **"import os" print(os.environ)**.

One of the first things you'll want to know is the current value for specific environment variables in your session. For that, you can use **os.environ.get()**. This command retrieves the value of an environment variable currently set on your system.

Key Topic

OpenSSL

OpenSSL is a commercial-grade, full-featured toolkit for the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols. It is also a general-purpose cryptography library. The ***OpenSSL*** libraries use environment variables to override the compiled-in default paths for various data. To avoid security risks, the environment is usually not consulted when the executable is set-user-ID or set-group-ID. CTLOG_FILE specifies the path to a certificate transparency log list.

The configuration file is called openssl.cnf by default and belongs in the same directory as openssl.exe by default. You can specify a different configuration file by using the OPENSSL_CONF environment variable, or you can specify alternative configurations within one configuration file.

The configuration file is a text file and comprises several sections, such as the ca section, which configures the certificate authority (CA). You can have several ca sections, each specifying a different configuration for a different CA, and switch

between them by changing the default_ca option. You can also override this choice from the command line, using the *-name* parameter. This capability is useful in development and testing, enabling you to try out different configurations.

The policy section specifies how closely the Distinguished Name in a certificate presented to SSL software must agree with the Distinguished Name in an installed certificate for the two certificates to be considered to match. The req section configures the **openssl req** command.

Packet Capture and Replay

A ***packet capture*** is a networking term for intercepting data packets that are crossing a specific point in a data network. Once a packet is captured in real time, it is stored for a period of time in a buffer or disk so that it can be analyzed and then either downloaded, archived, or discarded. Packets are captured and examined to help diagnose and solve network problems, provide forensic information, address cybersecurity, and enable research—in other words, when you're threat hunting. Replaying a packet capture allows organizations to test their intrusion detection rules, test their security operations center, and train analysts to discern what type of traffic is crossing their network so that they can learn to identify and respond. Tools that can be used to capture packets are tcpdump and tcpreplay on Linux and Wireshark on both Linux and Windows.

Tcpreplay

Tcpreplay is a suite of free open-source utilities for editing and replaying previously captured network traffic. Originally designed to replay malicious traffic patterns to intrusion detection/prevention systems, it has seen many evolutions, including capabilities to replay traffic to web servers. The basic operation of tcpreplay is to resend all packets from input files at the speed at which they were recorded, or a specified data rate, up to as fast as the hardware is capable. Traffic can be split between two interfaces, written to files, and filtered and edited in various ways, providing the means to test firewalls, network intrusion detection systems (NIDS), and other network devices.

NOTE For more details, see the Tcpreplay Manual at https://tcpreplay.appneta.com.

Tcpdump

Tcpdump is a data-network packet analyzer that runs under a command-line interface. It allows you to display TCP/IP and other packets being transmitted or received over a network to which the computer is attached. Distributed under the BSD license,

Tcpdump is free software. It is often used to help troubleshoot network issues; it is also a security tool used to collect packets between firewalls and the raw Internet.

Tcpdump can be used in a variety of cases. Because it's a command-line tool, it is ideal to run on remote servers or devices for which a GUI is not available; this way, you can collect data that can be analyzed later. It can also be launched in the background or as a scheduled job using tools like cron. To launch Tcpdump, you run **tcpdump -i** *interface* [*switches*].

Key Topic

Wireshark

Wireshark is one of the most widely used network protocol analyzers. It lets you see what's happening on your network at a microscopic level and is the de facto (and often de jure) standard across many commercial and nonprofit enterprises, government agencies, and educational institutions. Wireshark allows you to capture packets on wireless, Bluetooth, USB, and LAN network-attached cards, placing them into promiscuous mode. Wireshark is used to troubleshoot network problems, examine security problems, verify network applications, debug protocol implementations, and learn network protocols present on networks. It is available on Windows, Linux, and macOS; there are both command-line and GUI versions.

You can also save off-packet captures, load previously saved captures, and use the GUI to search and filter views, as shown in Figure 26-21. Wireshark decodes the various packets in various colors and provides the packet numbers, packet time, source and destination, and protocol.

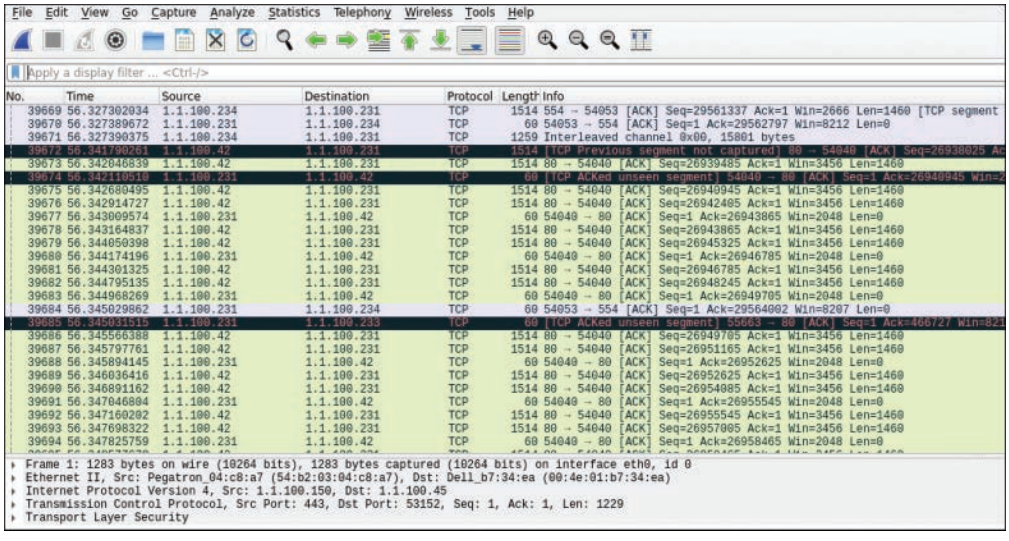


FIGURE 26-21 Wireshark GUI

Forensics

Key Topic

Digital ***forensics*** is the modern-day version of computer science forensics that deals with the recovery and investigation of material found in digital devices. It is most often used in cybercrime situations, including but not limited to, the attribution of bad actors, the identification of leaks within an organization, and the assessment of any damage that occurred during a breach. Using the data collected from electronic devices, digital forensic investigators can stop hackers and other cybercriminals from compromising an organization's digital infrastructure. They can also assist in recovering lost or stolen data, discover where a specific attack came from and trace it back to the source, and help create a detailed investigative report that can remedy any crime. The forensics field is spread across a number of subdivisions, all of which depend on the nature of the digital device that is the subject of the investigation, such as computer forensics, network forensics, forensic data analysis, and mobile device forensics.

Digital forensics and cybersecurity are connected in many ways and provide information to each other. When analyzing a breach, you may produce insights that could prevent future breaches and understand how particular threats make it easier for cybersecurity professionals and digital forensics investigators to establish a timeline and provide a direction for the investigation to follow.

Digital Forensics and Incident Response (DFIR) applies forensics to examine cases involving data breaches and malware, among other types of attacks. In many cases, digital forensics investigators have a background in computer sciences, which can help them develop the knowledge necessary to understand how virtual networks tick and work with one another. Most importantly, perhaps, they understand the vulnerabilities that exist within these systems and how they can be compromised.

dd

Key Topic

dd is a command-line utility for Linux operating systems whose primary purpose is to convert and copy files. On Linux, device drivers for hardware (such as hard disk drives) and special device files such as /dev/zero and /dev/random appear in the file system just like normal files. dd can also read from or write to these files, provided that function is implemented in their respective drivers. As a result, dd can be used for tasks such as backing up the boot sector of a hard drive and obtaining a fixed amount of random data. The dd program can also perform conversions on the data as it is copied, including byte order swapping and conversion to and from the ASCII and EBCDIC text encodings.

The command-line syntax of **dd** differs from many other Linux programs, in that it uses the syntax *option=value* for its command-line options, rather than the more standard *-option value* or *-option=value* formats. By default, **dd** reads from stdin and writes to stdout, but you can change this by using the **if** (input file) and **of** (output file) options.

Some practical examples of the **dd** command include the following:

- **To back up the entire hard disk:** To back up an entire copy of a hard disk to another hard disk connected to the same system, you execute the following **dd** command: **$ dd if = /dev/sda of = /dev/sdb**. In this **dd** command example, the Linux device name of the source hard disk is /dev/hda, and the device name of the target hard disk is /dev/hdb.
- **To create an image of a hard disk:** Instead of making a backup of the hard disk, you can create an image file of the hard disk and save it in other storage devices. There are many advantages of backing up your data to a disk image, one being the ease of use. This method is typically faster than other types of backups, enabling you to quickly restore data following an unexpected catastrophe. The following command creates the image of a hard disk /dev/hda: **# dd if = /dev/hda of = ~/hdadisk.img**. This image can be used to preserve forensic evidence of a computer system that was attacked or used to exploit other systems.

Key Topic

Memdump

Memdump is a memory dump, which is the process of taking all information content in RAM and writing it to a storage drive. Developers commonly use memory dumps to gather diagnostic information at the time of a crash to help them troubleshoot issues and learn more about the event. Information yielded by the memory dump can help developers fix errors in operating systems and other programs of all kinds. Some computer errors are unrecoverable because they require a reboot to regain functionality, but the information stored in RAM at the time of a crash contains the code that produced the error. Memory dumps save data that might otherwise be lost to RAM's volatile nature or overwriting.

You might have seen a memory dump as the blue screen of death on a Microsoft operating system. The errors display some basic suggestions, information, and a faulting module while the percentage of the memory written to storage counts up. Following reboot, the memory dump can be sent to a memory image server or to Microsoft for analysis to help the company fix the issue in updates and learn about usage.

Because these dumps can include anything in the computer's active RAM, some users have privacy concerns. Furthermore, because the dumps are stored on the drive, they can also present security risks. Modern hackers are aware of these issues and initially look for memory dumps on systems. After obtaining them, they can potentially find cleartext passwords or decryption keys that normally would not be easily accessible. Certain operating systems allow you to restrict collected information about memory dumps, and some make it possible to turn off memory dumps entirely.

Key Topic

WinHex

WinHex, made by X-Ways Software Technology AG of Germany, is a powerful application; an advanced hex editor; a tool for data analysis, editing, and recovery; a data editing and wiping tool; and a forensics tool used for evidence gathering. At its most basic level, WinHex is a hex editor. That is, like most hex editors, it displays three columns: an address, a 16-byte hex display, and a 16-character text display.

The data viewer can be extensively configured. For example, by clicking the up, down, right, and left arrows on the toolbar, you can add lines, remove lines, add columns, and remove columns from the data display. You can also view hex only, text only, or both by clicking checkboxes in the View menu. The more general options let you set the colors and font, and clicking the Offset column toggles between decimal and hex address values. Like other hex editors, WinHex can open files as editable or as read-only. If you open your hard disk as editable and make a change, you could destroy the drive. When you edit a file, the data change is stored in a temporary file until saved, at which time your changes are committed. There is also an in-place Edit mode in which all changes are made directly in real time (the default when editing RAM). You can choose the Edit mode from the Open File dialog. When you open an entire disk or partition, the default mode is Edit, as shown in Figure 26-22.

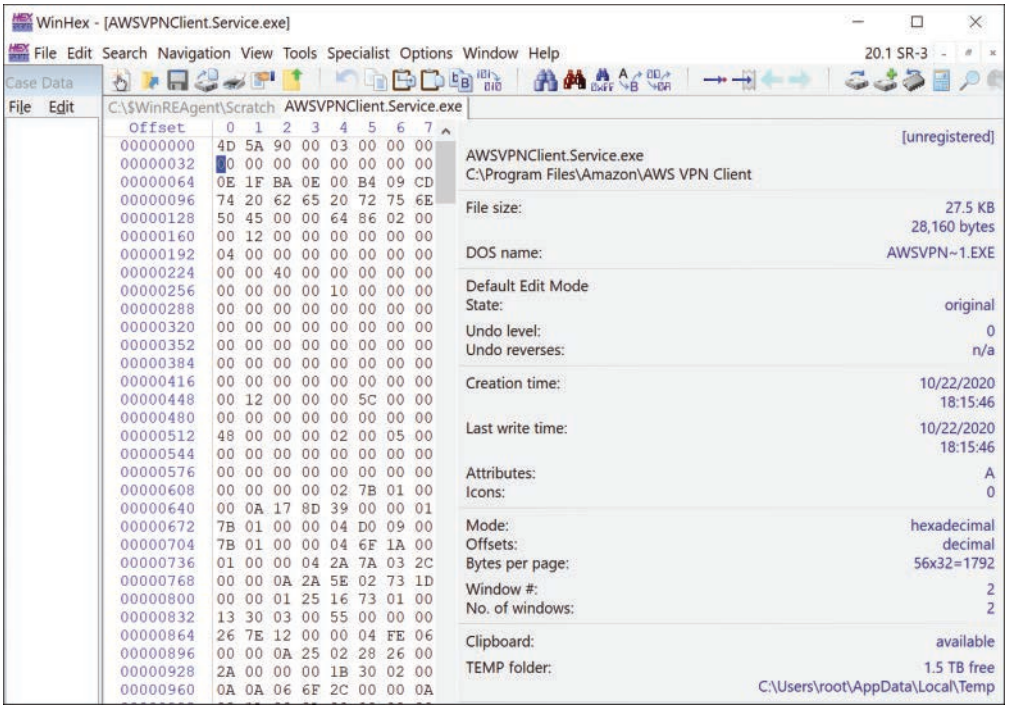


FIGURE 26-22 WinHex in View Mode

FTK Imager

FTK Imager is a data preview and imaging tool that lets you quickly assess electronic evidence to determine whether further analysis with a forensic tool such as AccessData's Forensic Toolkit (FTK) is warranted. FTK Imager can create perfect copies or forensic images of computer data without making changes to the original evidence. The forensic image is identical in every way to the original, including file slack and unallocated space or drive free space. Using this tool allows you to store the original media away, safe from harm while the investigation proceeds using the image. FTK can generate hash reports for regular files and disk images to use as a benchmark to prove the integrity of your case evidence. When a full drive is imaged, a hash generated by FTK Imager can be used to verify that the image hash and the drive hash match after the image is created and that the image has remained unchanged since acquisition.

Autopsy

Autopsy is a digital forensics platform and graphical interface to The Sleuth Kit and other digital forensics tools. It is used by law enforcement, military, and corporate examiners to investigate what happened on a computer. You can even use it to recover photos from your camera's memory card.

Autopsy was designed to be an end-to-end platform with modules that come with it out of the box and others that are available from third parties. Some of the modules provide these capabilities:

- **Timeline analysis:** Use advanced graphical event viewing interface.
- **Hash filtering:** Flag known bad files and ignore known good ones.
- **Keyword search:** Search indexed keywords to find files that mention relevant terms.
- **Web artifacts:** Extract history, bookmarks, and cookies from web browsers like Firefox, Chrome, and Microsoft Edge.
- **Data carving:** Recover deleted files from unallocated space using PhotoRec.
- **Multimedia:** Extract EXIF from pictures and watch videos.
- **Indicators of compromise:** Scan a computer using STIX.

Exploitation Frameworks

Exploitation frameworks are software packages that contain reliable exploit modules and other hacker technique tools such as agents used for successful repositioning. Exploitation frameworks allow you to use different exploit payloads and other

unique options to obfuscate shell code and network traffic to avoid detection. Many of the currently supported exploitation frameworks have significant support models in place to help organizations perform their pen testing and red team attack analysis internally, usually in conjunction, prior to a third-party analysis. Three of the most notable frameworks are Metasploit, Core Impact, and Immunity Canvas; there also are a number of less famous frameworks.

Organizations of all sizes are faced with the challenges of maintaining a successful patch management program. In many cases, vulnerability scans and software updates are performed only on a monthly basis. The lack of visibility into the network and systems in between active scans can result in an increased risk to the organization.

This point-in-time method of scanning and updating can result in systems being missed if they are not on the network or available during the scan window. One vulnerability is often the only necessary piece needed to gain a foothold in an environment. As an example, a network could be compromised due to a vulnerability found in out-of-date office productivity software, PDF viewers, or browsers. Exploitation framework tools contain capabilities to detect and exploit these vulnerabilities.

The vendors of these software packages are continually adding exploits to their platforms. Internal security teams and malicious actors alike can use the same tools to detect and exploit vulnerabilities. Because some of the software exploitation tools are free, the bar of entry is minimal and can open up organizations to easily to perform attacks.



Password Crackers

A ***password cracker*** is an application program that is used to identify an unknown or forgotten password to a computer or network resources. Hackers can also use it to obtain unauthorized access to resources. Password crackers use two primary methods to identify correct passwords: brute-force and dictionary searches. When a password cracker uses a brute-force method, it runs through combinations of characters within a predetermined length until it finds the combination accepted by the computer system.

When conducting a dictionary search, a password cracker searches each word in the dictionary for the correct password. Password dictionaries exist for a variety of topics and combinations of topics, including politics, movies, and music groups. Some password cracker programs search for hybrids of dictionary entries and numbers. For example, a password cracker may search for ants01, ants02, ants03, and so on. This capability can be helpful where users have been advised to include a number in their password.

A password cracker may also be able to identify encrypted passwords. After retrieving the password from the computer's memory, the program may be able to decrypt it. Or, by using the same algorithm as the system program, the password cracker creates an encrypted version of the password that matches the original. Password crackers can also be used by organizations to identify weak and reused passwords in an audit.

Password cracking refers to the process of extracting passwords from the associated password hash. This can be accomplished in a few different ways:

- **Dictionary attack:** Most people use weak and common passwords. Taking a list of words and adding a few permutations—like substituting $ for *s*—enables a password cracker to learn a lot of passwords very quickly.
- **Brute-force guessing attack:** There are only so many potential passwords of a given length. While slow, a brute-force attack (trying all possible password combinations) guarantees that an attacker will crack the password eventually.
- **Hybrid attack:** A hybrid attack mixes these two techniques. It starts by checking to see if a password can be cracked using a dictionary attack and then moves on to a brute-force attack if it is unsuccessful.

Most password-cracking or password finder tools enable a hacker to perform any of these types of attacks. Some of the most commonly used password-cracking tools include the following:

- **Hashcat:** One of the most popular and widely used password crackers in existence. It is available on every operating system and supports over 300 different types of hashes. Hashcat enables highly parallelized password cracking with the ability to crack multiple different passwords on multiple different devices at the same time and the ability to support a distributed hash-cracking system via overlays.
- **John the Ripper:** A well-known free open-source password-cracking tool for Linux and macOS. A Windows version is also available.

John the Ripper offers password cracking for a variety of different password types. It goes beyond OS passwords to include common web apps (like WordPress), compressed archives, document files (Microsoft Office files, PDFs, and so on), and more.

- **THC Hydra:** A rapid password-cracking tool that attempts to determine user credentials by performing a brute-force password guessing attack. It is available for Windows, Linux, Free BSD, Solaris, and macOS. THC Hydra is extensible with the ability to easily install new modules. It also supports a large number of network protocols.

Data Sanitization

Data sanitization is the process of irreversibly removing or destroying data stored on a memory device (hard drives, flash memory/SSDs, mobile devices, CDs, DVDs, and so on) or in hard copy form. It is important to use the proper technique to ensure that all data is purged. Data sanitization is the process of deliberately, permanently, and irreversibly removing or destroying the data stored on a memory device to make it unrecoverable. A device that has been sanitized has no usable residual data, and even with the assistance of advanced forensic tools, the data will not ever be recovered.

There are three methods to achieve data sanitization: physical destruction, cryptographic erasure, and data erasure. However, the downside of these techniques is that they damage the storage media and do not allow it to be sold or reused. They are complex and expensive to carry out and are also harmful to the environment. The data erasure process uses software to write random 0s and 1s on every sector of the storage equipment, ensuring no previous data is retained.

This is a reliable form of sanitization because it validates that 100 percent of the data was replaced at the byte level. It is also possible to generate auditable reports that prove data has been successfully sanitized. The advantage of this method compared to physical destruction is that it does not destroy the device and allows it to be sold or reused.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 26-8 lists a reference of these key topics and the page number on which each is found.

Table 26-8 Key Topics for Chapter 26

Key Topic Element	Description	Page Number
Section	tracert/traceroute	707
Section	nslookup/dig	709
Section	ipconfig/ifconfig	710
Section	nmap	711

Key Topic Element	Description	Page Number
Section	ping/pathping	714
Section	hping	717
Section	netstat	718
Section	netcat	720
Section	IP Scanners	721
Section	arp	721
Section	route	723
Section	curl	724
Section	theHarvester	725
Section	sn1per	726
Section	scanless	727
Section	dnsenum	728
Section	Nessus	730
Section	Cuckoo	731
Section	head	733
Section	tail	734
Section	cat	734
Table 26-6	grep command options	736
Section	chmod	736
Table 26-7	logger command options	738
Section	Shell and Script Environments	738
Section	SSH	739
Section	PowerShell	740
Section	Python	741
Section	OpenSSL	741
Section	Packet Capture and Replay	742
Section	Tcpreplay	742
Section	Tcpdump	742
Section	Wireshark	743
Section	Forensics	744
Section	dd	744

Key Topic Element	Description	Page Number
Section	Memdump	745
Section	WinHex	746
Section	FTK Imager	747
Section	Autopsy	747
Section	Exploitation Frameworks	747
Section	Password Crackers	748
Section	Data Sanitization	750

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

network reconnaissance, tracert/traceroute, nslookup, dig, ipconfig, ifconfig, nmap, ping, pathping, hping, netstat, netcat, IP scanners, route, curl, sn1per, scanless, dnsenum, Nessus, Cuckoo, head, tail, cat, grep, chmod, logger, Shell variables, SSH, PowerShell, Python, OpenSSL, packet capture, Tcpreplay, Tcpdump, Wireshark, forensics, dd, Memdump, WinHex, FTK Imager, Autopsy, exploitation frameworks, password cracker, data sanitization

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What IP tool on Windows and Linux measures transit delays between packets across a network?
2. What open-source software allows you to take a suspicious file, isolate it, and run tests to provide a report on its behavior?
3. What IP level tool is mainly used to verify connectivity to other hosts and uses ICMP?
4. Which IP tool supports TCP, UDP, ICMP, and RAW IP protocols; has the ability to send files and perform firewall testing; and has many advanced features including operating system fingerprinting?
5. The curl tool can be used to transfer data from host to host by using which protocol?

This page intentionally left blank



This chapter covers the following topics related to Objective 4.2 (Summarize the importance of policies, processes, and procedures for incident response) of the CompTIA Security+ SY0-601 certification exam:

- Incident response plans
- Incident response process
  - Preparation
  - Identification
  - Containment
  - Eradication
  - Recovery
  - Lessons learned
- Exercises
  - Tabletop
  - Walkthroughs
  - Simulations
- Attack frameworks
  - MITRE ATT&CK
  - The Diamond Model of Intrusion Analysis
  - Cyber Kill Chain
- Stakeholder management
- Communication plan
- Disaster recovery plan
- Business continuity plan
- Continuity of operations planning (COOP)
- Incident response team
- Retention policies

Summarizing the Importance of Policies, Processes, and Procedures for Incident Response

A recent report recognizes that nearly three-quarters of organizations don't have a consistent enterprise wide cybersecurity incident response (IR) plan. Organizations with IR teams and regular testing have an average data breach costing $2 million less than those with no IR team or plan in place. Today's enterprises are global, as are their footprints. Having applications and resources in the cloud and remote workforces has increased the potential for cyber threats. The first and most important step in the incident response lifecycle is preparation. Preparation ahead of an incident is what will allow you to respond more quickly and effectively in the midst of an incident.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 27-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 27-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Incident Response Plans	1–2
Incident Response Process	3–4
Exercises	5
Attack Frameworks	6–7
Stakeholder Management	8

Foundation Topics Section	Questions
Communication Plan	9
Disaster Recovery Plan	10–11
Business Continuity Plan	12
Continuity of Operations Planning (COOP)	13
Incidence Response Team	14
Retention Policies	15

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which high-level document is a step-by-step procedure that should be created as part of an incident response plan that can target specific incident handling like malware and ransomware?
   a. Playbooks
   b. Play stations
   c. Lessons learned after action report
   d. Security incident field report

2. Which legal portion of an incident response plan requires notification or disclosure within 72 hours of discovery of a data incident?
   a. DR Disclosure Federal Law
   b. NIST Disclosure Law
   c. New Jersey Privacy Law
   d. GDPR

3. Which incident response plan item will provide an understanding of the severity of an incident so that it can be prioritized quickly and correctly?
   a. Disaster recovery report
   b. Incident response report
   c. Triage matrix
   d. Threat matrix

4. Which phase of the incident response process should be performed within two weeks of the end of an incident?
   a. Identification
   b. Preparation
   c. Lessons learned
   d. Remediation

5. Which exercise simulates a real-life scenario of an incident response plan and is used to the test and highlight areas where your team excels and areas that need to be addressed?
   a. Tabletop
   b. Containment exercises
   c. Recovery process
   d. Cyber kill chain

6. The Diamond Model places the basic components of malicious activity at one of the four points on a diamond shape. What are the four points? (Choose two)
   a. Malware and infection vectors
   b. Personas and biometrics
   c. Adversary and infrastructure
   d. Capability and victim

7. The ATT&CK Framework has 11 tactics and hundreds of techniques. Which tactic describes the way an adversary implements a technique?
   a. Collection
   b. Procedures
   c. Privilege escalation
   d. Impact

8. Which of the following is one of the five key stakeholders of the incident response team?
   a. Security Operations
   b. Security Guards
   c. Public Library
   d. Legal

9. In a communication plan, escalating communication information on a regular schedule or timeline is important. What is the appropriate frequency of this communication?
   a. Once an hour
   b. Once every six hours
   c. As key information is available
   d. As every item is uncovered

10. Which of the following is a formal document that contains details on how to respond to a cyber attacks and unplanned incidents?
   a. Incident response model
   b. Disaster continuity plan
   c. Disaster recovery plan
   d. Containment process

11. Not implementing a disaster recovery plan properly can lead to which of the following?
   a. Satisfied customers
   b. Brand awareness
   c. Lost revenue
   d. Faster recovery

12. What is one thing that a BCP plan contains that a DR plan does not?
   a. A standby data center
   b. Continuity of a DRP in conjunction with a BCP
   c. A continuity plan for the entire organization
   d. A disaster recovery model

13. Which of the following ensures the restoration of organizational functions in the shortest possible time?
   a. COOP
   b. MITRE ATT&CK
   c. Diamond Model
   d. Cyber kill chain

14. What are some of the incidents that an incident response team might be prepared for and respond to? (Select all that apply.)
    a. Attackers gaining access to the web server
    b. Hackers obtaining passwords from executives
    c. A nasty computer virus that the antivirus contained
    d. A power outage in the data center

15. NIST SP 800-53 requires that all federal agencies retain data for how many years?
    a. Ten years on magnetic media or 20 years on paper
    b. Three years on magnetic media
    c. Seven years on magnetic media
    d. Seven years on magnetic media and 10 years on paper

Foundation Topics

Incident Response Plans

In the simplest of terms, a cybersecurity ***incident response plan*** (or IR plan) is a set of instructions designed to help companies prepare for, detect, respond to, and recover from network security incidents. An incident response plan ensures that in the event of a security breach, the right personnel and procedures are in place to effectively deal with a threat. Having an incident response plan in place ensures that a structured investigation can take place to provide a targeted response to contain and remediate the threat.

An organization's failure to have or implement an incident response plan can have serious legal repercussions. To effectively deal with a cybersecurity incident, your company may need a team that specializes in incident response. Some organizations call this team the computer security incident response team (CSIRT); there are other permutations of that acronym like security incident response team (SIRT) or computer incident response team (CIRT). The mission of this team is the same no matter what you call it—to enact the company's established incident response plan when a cybersecurity incident occurs.

If you work in data security, you deal with security incidents on a day-to-day basis. Occasionally, a minor security issue turns out to be a real-life panic situation. When an incident occurs, will everyone on the team know what to do? Will CSIRT members know their role and responsibilities and follow the approved plan?

Simply having an IR plan is not enough; the CSIRT team must have the skills and experience to deal with a potentially high-stress situation like this. The team needs digital forensics experts, malware analysts, incident managers, and security operations center (SOC) analysts who are all heavily involved and actively dealing with the situation. This involves making key decisions, conducting an in-depth investigation, providing feedback to key stakeholders, and ultimately giving assurances to senior management that the situation is under control.

This activity often takes place in a time crunch. Data breach notification laws are becoming more common: the General Data Protection Regulation (GDPR) in the European Union, for instance, requires that companies report data security incidents within 72 hours of discovery.

Having an incident response plan is imperative. The first step is identifying and having the right people with the right skill sets and experience available and ready to respond. You should regularly test and update your incident response plan. Everyone

who is part of the plan should understand their role and the role of others to help reduce confusion during a real event.

Key Topic Incident Response Process

The incident response process is made up of four key elements that can be developed as a company's security posture measures. There are important considerations to be made when building an incident response plan. First and foremost, backing from senior management is paramount. Occasionally, people use the incident response process and IR plan interchangeably. It's important to note that the process is a roadmap to developing a specific plan for each organization; each is as different as the organization.

Building an incident response plan should not be a box-checking exercise or something handed off to inexperienced employees. If senior management does not support this process, there is a risk of its becoming filed away, incomplete and useless when needed. Senior leadership should be outlining critical processes, systems, and resources important to business continuity. Part of plan development includes defining the key stakeholders and obtaining contact details for key individuals and teams inside and outside of business hours; this information needs to be added to the plan.

Figure 27-1 shows the NIST incident response lifecycle. The incident response process is a business process that enables you to remain in business. The list that follows describes the four phases in more detail.

NOTE For more details on NIST incident handling, see https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf.

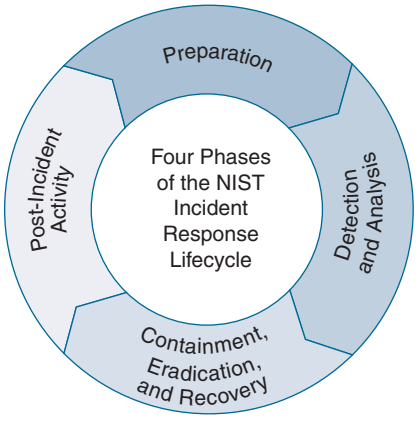


FIGURE 27-1 The Four Phases of the Incident Response Lifecycle

Phase 1. **Preparation:** The quality of the response to an incident largely depends on incident response preparedness. In this phase, all components needed to effectively respond to an incident are identified, acquired, or created.

Phase 2. **Detection and analysis:** This phase primarily focuses on detection and discovery of indicators of compromise (IOCs). Here, having an incident reporting policy and procedure in place is critical to training.

Phase 3. **Containment, eradication, and recovery:** After stopping the problem from getting worse or spreading, you limit the damage and then regain control of your network and systems. Finally, recovery starts with restoration of systems to normal operations.

Phase 4. **Post-incident activity:** This phase centers on lessons learned to improve the incident response capability and prevent the incident from happening again.

A playbook provides manual orchestration of incident response. For example, specific incidents and threats have their own playbooks. As a result, the response that an organization takes is formalized in a step-by-step procedure.

By contrast, IT operations or security operations centers (SOCs) use runbooks as a reference for routine procedures that administrators perform both as standard mode of operations and during emergencies such as an IR incident.

A security operations center is similar to a network operations center in that it is staffed 24/7 by expert staff. In an SOC, security experts monitor, alert, respond to, and triage incidents. They also act as the front line for all security-related incidents.

Preparation

The incident response plan should have an owner, and that owner is responsible for sending out communications, assigning tasks, and establishing the appropriate actions that should be taken. Also, your organization should consider who needs to be included in any incident communication and how much detail is required depending on the audience. Tasks assigned to security teams need to be precise and technical, whereas updates to the company board (executives) need to be clear and free of any technical terms. You should develop playbooks that provide guidance to the SOC when triaging an incident; they should give clear instructions on how to prioritize an incident and when and how incidents should be escalated. The playbooks should be high level and focused on specific areas such as malware, insider threats, unauthorized access, ransomware, and phishing. The playbooks and procedures should be tested with the people and teams who will be using them. Tabletop exercises are an excellent way to solidify the knowledge and see whether any improvements can be made. Playbooks and tabletop exercises are described in more detail shortly.

TIP A risk matrix displays the probability or likelihood versus the consequences of risk. Table 27-2 provides a sample matrix.

Table 27-2 Risk Matrix

Likelihood	Consequences				
	Not Significant	**Minor**	**Moderate**	**Major**	**Severe**
Very Likely	Medium	High	High	Very High	Very High
Likely	Medium	High	High	Very High	Very High
Possible	Low	Medium	High	High	Very High
Unlikely	Low	Low	Medium	Medium	High
Rare	Low	Low	Low	Low	Medium

Identification

According to an old saying, you can only successfully defend what you see. In this case, you can only remove a security threat when you know the size and scope of the incident. Step one begins with identifying "patient zero," the initially compromised device. The goal here is to understand the root cause of the compromise; however, you should not just focus on the one device. Could the threat have spread and moved laterally, or is there another potential initially compromised device?

Actual identification of an incident comes from gathering useful indicators of compromise. Investigators should look to identify any unique IOCs that can be used to search across your network for further evidence of compromise. If the incident relates to a malware infection, then ask the following questions: What network traffic and connections does the malware generate? Does the malware connect to any specific IP address or domain? What files were downloaded? What running processes are created? What files are created in memory or on disk? Have any unique registry keys been created? This data is used to search for further evidence of compromise and identify any other infected machines in your network.

Containment

Once the scope of an incident has been successfully identified, the containment process can then begin. This is where the compromised devices within the network are isolated from the rest of the network to stop the spread. Short-term containment may be used to isolate a device that is being targeted by attack traffic. Long-term containment may be necessary when a deep-dive analysis is required, which can be

time consuming. This process may involve taking a forensic image of the device and conducting detailed forensic analysis; the analysis may generate further IOCs and identification of the source, and may need to be revisited.

Eradication

Many organizations don't understand the risks associated with improper eradication and how to actually fix a malware infection or breached system. Once the incident is successfully contained, the eradication of the threat can begin. This process varies depending on what devices have been compromised and what caused the compromise. If possible, you should perform a complete wipe/reimage of the affected system. This step, of course, requires you to have good backups in place and the ability to establish the initial date/time of infection, rolling back to just before that. Besides reimaging systems, you should make sure you use the latest patches on all devices, disarm malware, disable compromised accounts, and change passwords. These are just a few examples of what may be required in the eradication phase of an incident.

Regardless of how you choose to eradicate an infection, you need to have a plan for immediate increased monitoring of any affected systems for some period of time after the eradication process, usually 30 days or longer. This plan is important to make sure the steps you took to fix the issue were effective and there were no lingering malware, rootkits, backdoors, or additional compromised accounts. This should be a daily task where you review event logs from both affected host and Active Directory logs for account usage.

Recovery

The IR team brings affected production systems back online carefully to ensure another incident doesn't take place. Important decisions at the recovery stage are from which time and date to restore operations, how to test and verify that affected systems are back to normal, and how long to monitor the systems to ensure activity is back to normal. The goal of the recovery phase of an incident is to restore normal service to the business. If clean backups are available, then they can be used to restore service. Alternatively, any compromised device will need rebuilding to ensure a clean recovery. Additional monitoring of affected devices may need to be implemented.

Lessons Learned

The lessons learned phase should be performed no later than two weeks from the end of the incident. A Post Incident Review (PIR) meeting is used to ensure information is fresh on the team's mind. The main purpose of this phase is to update documentation that could not be prepared during the response process and investigate

the incident further to identify its full scope, how it was contained and eradicated, and what was done to recover the attacked systems. After the threat has been fully remediated, the next step involves answering this question: How do we stop this from happening again? The (PIR) is the platform used to discuss what went well during the incident and what processes and procedures can be improved. At this stage, the incident response plan is refined based on the outcome of the PIR, and procedures and playbooks are amended to reflect any agreed changes. Figure 27-2 illustrates the stages of the incident response process.

Key Topic

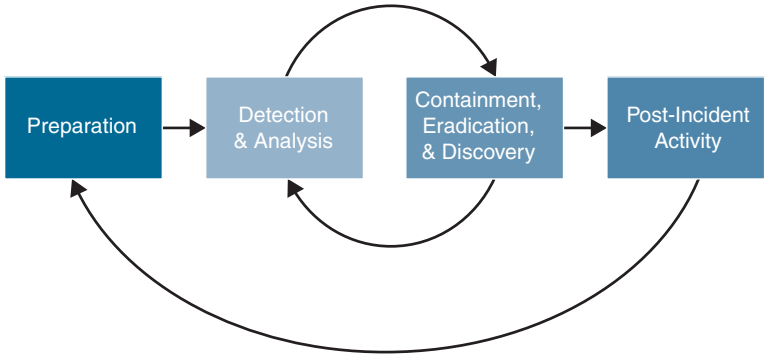


FIGURE 27-2 Incident Response Process

Exercises

IR exercises are an excellent way to solidify your organization's knowledge and see if any improvements can be made. From these exercises, you want to ensure your organization's increased awareness and understanding of threats, be able to evaluate your overall incident preparedness, and identify deficiencies in your IR plan, including technical, planning, procedural, and resource gaps.

Key Topic

Tabletop

A ***tabletop exercise*** begins with a security incident preparedness activity, a plan to test your IR plan. It takes participants through the process of dealing with a simulated incident scenario and provides a hands-on training exercise for participants who can then highlight flaws in incident response planning. Usually, a trained expert facilitates the discussion through multiple scenarios to determine the team's readiness or whether there are potential gaps. The output of this exercise is used to drive enhancements to the organization's plan and approach to identifying, analyzing, and resolving incidents and how they could be prevented in the future.

A tabletop exercise is used to validate and improve an organization's IR plan. Real-life scenarios are used to put the response plan to the test, highlighting areas where your team excels and areas to be addressed. The tabletop exercise also ensures that everyone on your team knows their roles and responsibilities in the event of an attack. The tabletop exercise aligns everyone's understanding of the process, helps assign the roles of each member, and provides participants with knowledge through hands-on experience. The exercise begins with the incident response plan (IRP) in a classroom-type setting and gauges team performance against the following questions: What happens when you encounter a breach? Who does what, when, how, and why? What roles will HR, legal, IT, corporate communication, and company officers play? Who is assigned to spearhead the effort, and what specific authority will they have? What resources are available to them and when they need them?

Key Topic

Walkthroughs

The facilitator develops the input required to meet the predetermined goal of the exercise, based on the preset specifications about the environment and goals set by leadership. Walkthroughs take place prior to tabletop exercises. Some of these steps depend on the industry, regulations, and their current security posture. The facilitator develops the test input to meet a specific level of management, as well as the appropriate technical personnel and all others who should make the appropriate decisions. In this discussion-based exercise, you walk through the steps of the plan. Typically, a project kickoff call/meeting occurs between the facilitator and the team to assign roles and responsibilities, update the IR plan, and identify any other participants needed. A typical walkthrough exercise flow consists of three essential elements: inputs, process, and outputs.

Key Topic

Simulations

Security incident response simulations (SIRS) are internal events that provide a structured exercise to practice your team's IR procedures and plan in a simulated realistic scenario, where components such as deadlines and external injects increase the realism of the event. Injects are scenario-based actions that help add a realistic element to the event. SIRS events are about preparing team members with realistic simulations and helping them improve response capabilities. The value in these exercises is what is obtained, retained, and fed back into your plan, starting with

- Validating the team and the plan's readiness
- Identifying and documenting deficiencies to facilitate constructive feedback
- Helping develop team member confidence
- Providing evidence of compliance, depending on your industry

The value and benefit derived from team members participating in SIRS include increases in the organization's effectiveness during stressful live events.

Attack Frameworks

There are multiple frameworks for incident response, including one from NIST and one from SANS. These are the two dominant institutes whose incident response steps have become the industry standard. Similarly, a number of attack frameworks are available for use in attack scenarios. The benefit of a framework is that it puts everyone on the same page, with a common language, and a common place to level set and start so that everyone is speaking the same terms. Imagine if three people in the same organization used three different frameworks to communicate about an attack or incident. Besides this scenario being counterproductive, the company would waste valuable time trying to identify and recover from an incident.

Key Topic

MITRE ATT&CK

A ***MITRE ATT&CK*** is a globally accessible knowledge base of adversary tactics, techniques, and procedures (TTPs) based on real-world observations of cybersecurity threats. They're displayed in multiple matrixes that are arranged by attack stages, from initial system access to data theft or machine control.

The aim of the framework is to improve post-compromise detection of adversaries in enterprises by illustrating the actions attackers may have taken. How did attackers get in? How are they moving around? The knowledge base is designed to help answer these questions while contributing to the awareness of an organization's security posture at the perimeter and beyond. Organizations can use the framework to identify holes in defenses and prioritize remediation of them based on risk.

Threat hunters leverage the ATT&CK framework to look for specific tactics, techniques, and procedures that adversaries may use in conjunction with other methods. The framework can be useful for gauging an environment's level of visibility against targeted attacks with the existing tools deployed across an organization's endpoints and perimeter.

The enterprise ATT&CK framework consists of 11 tactics. You might consider tactics the "why" part of the ATT&CK equation. What objective did attackers want to achieve with the compromise?

- Initial access
- Execution
- Persistence
- Privilege escalation

- Defense evasion
- Credential access
- Discovery
- Lateral movement
- Collection
- Exfiltration
- Impact

Each tactic contains an array of techniques that have been observed being used by threat actor groups in compromises or by malware. You can think of tactics as the "how" part of the ATT&CK framework. How are attackers escalating privileges or exfiltrating data?

Although there are only 11 tactics in the enterprise ATT&CK framework, there are scores of techniques, too many to list individually here, but as of this writing, there are 291. They're perhaps best researched and visualized via MITRE's ATT&CK Navigator, an open-source web application that allows basic navigation and annotation. Techniques are referenced in ATT&CK as T*xxxx*, where, for example, spear phishing is T1192 and remote access tools are T1219. Each technique contains contextual information, like the permissions required, what platform the technique is commonly seen on, and how to detect commands and processes they're used in.

Procedures describe the way adversaries implement a technique. A procedure concerns the way the instance was used. It also can be useful for understanding exactly how a technique is used and for replication of an incident with adversary emulation as well for specifics on how to detect the instance in use.

Although the framework has been around for years, it is being adopted by more organizations, the government, and end users to share threat intelligence. While there are other ways to share threat intelligence, the ATT&CK framework provides a common language that's standardized and is globally accessible.

NOTE To learn more about the ATT&CK framework, visit the Mitre site at https://attack.mitre.org/.

Key Topic

The Diamond Model of Intrusion Analysis

In the cybersecurity and threat intelligence industries, several approaches are used to analyze and track the characteristics of cyber intrusions by advanced threat actors. The ***Diamond Model of Intrusion Analysis*** emphasizes the relationships and

characteristics of four basic components: the adversary, capabilities, infrastructure, and victims (see Figure 27-3). The main axiom of this model states, "For every intrusion event, there exists an adversary taking a step toward an intended goal by using a capability over infrastructure against a victim to produce a result." This means that an intrusion event is defined as how the attacker demonstrates and uses certain capabilities and techniques over infrastructure against a target.

Key Topic

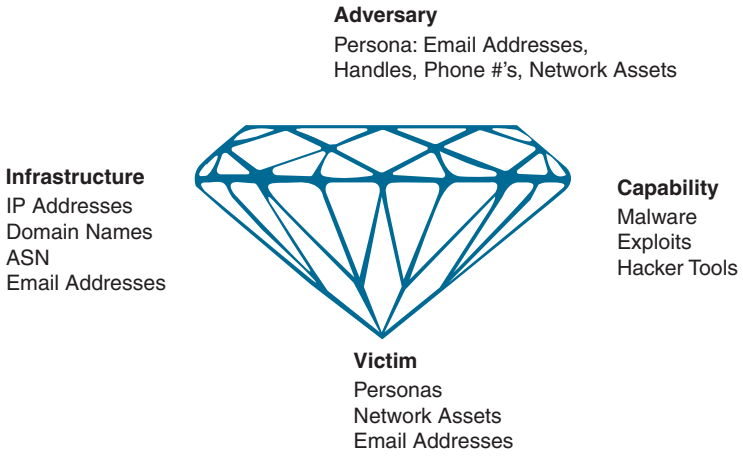


FIGURE 27-3 Diamond Model of Intrusion Analysis

The Diamond Model identifies adversaries with developing capabilities and techniques that are unique to that group. The method's context directly translates to the capability edge of that model. It can become obvious that an adversary uses distinct malware and attack vectors as part of its capabilities and TTPs.

Figure 27-3 shows specifically what an adversary might look like or what information may be gleaned from the intrusion, whereas infrastructure shows what was used, capability shows the method, and victim shows the activity that allowed the attacker in. By identifying events and linking them into activity threads, an analyst gains information regarding what occurred during an attack. By looking at the gaps in their knowledge (such as missing features), the analyst identifies where further information is needed.

MITRE ATT&CK and the Diamond Model are both very focused on the capabilities of attackers. The capabilities are the most unchanging attributes about a particular threat actor. The Diamond Model also discusses infrastructure, which can be used to link different attack campaigns to a particular adversary. However, this infrastructure can also be easily replaced and can present false positives, due to the fact that multiple distinct actors can use the same infrastructure, making it a less-than-ideal tool for attribution. However, an attacker's capabilities, on the other hand,

require significant investment to build and are likely more unchanging. After development of a custom piece of malware, a threat actor is likely to use it across multiple campaigns, providing a more reliable basis for attribution.

Cyber Kill Chain

The ***cyber kill chain*** is a series of eight steps that trace stages of a cyberattack from the early reconnaissance stages to the exfiltration of data.

The cyber kill chain enables you to understand and combat ransomware, security breaches, and advanced persistent threats (APTs). Lockheed Martin derived the cyber kill chain framework from a military model—originally established to identify, prepare to attack, engage, and destroy the target. Since its inception, the kill chain has evolved to better anticipate and recognize insider threats, ransomware, social engineering, and innovative and advanced types of attacks.

Key Topic

Each stage is related to a certain type of activity in a cyber attack, regardless of whether it's an internal or external attack:

1. **Reconnaissance:** This is the observation stage, where attackers typically assess the situation from the outside in order to identify both targets and tactics required for the attack.
2. **Intrusion:** This is based on what the attackers discovered in the reconnaissance phase and how they're able to get into your systems, which is often by leveraging malware or security vulnerabilities.
3. **Exploitation:** This is the act of exploiting vulnerabilities and delivering malicious code onto the system to get a better foothold.
4. **Privilege Escalation:** Attackers often need more privileges on a system to get access to more data and permissions: for this, they need to escalate their privileges, often to an Admin.
5. **Lateral Movement:** After they are in the system, attackers can move laterally to other systems and accounts to gain more leverage or control, whether that's higher permissions, more data, or greater access to systems.
6. **Obfuscation/Anti-forensics:** To successfully pull off a cyber attack, attackers need to cover their tracks, and in this stage, they often lay false trails, compromise data, and clear logs to confuse and/or slow down any forensics team.
7. **Denial of Service (DoS):** This is the disruption of normal access for users and systems to stop the attack from being monitored, tracked, or blocked.
8. **Exfiltration:** This is the extraction stage: getting data out of the compromised system.

Each phase of the cyber kill chain is an opportunity to stop a cyberattack in progress. With the right tools to detect and recognize the behavior of each stage, you are able to better defend against a systems or data breach.

NOTE The MITRE ATT&CK may be similar to a cyber kill chain but focuses more on the nuances of different attack techniques to understand and defend against an attacker. The Diamond Model places the basic components of malicious activity at one of four points of a diamond shape: adversary, infrastructure, capability, and victim. Lockheed Martin developed the cyber kill chain framework to help defend its networks based on the chain of actions that an attacker takes from beginning to end. This model has since been adopted by the security industry.

Key Topic

Stakeholder Management

Managing the expectations of leaders of your organization is one of the most important aspects of a complete incident response plan. Leaders of an organization must communicate to owners, stockholders, and other leadership within the organization. Ensuring your plan includes a regular interval in which you update management/leadership is essential to setting expectations and providing transparency in the process; this all should be addressed in the IR plan and be part of the tabletop exercise.

The use of a communication template is critical to appropriately word the message to stakeholders. Everyone on your IR team is a stakeholder, and communication and management of activities are important as well. There are five key stakeholders for any IR team, including IT Services, Security Management, Legal, Human Resources, and Public Relations. Another consideration—and one that requires management approval—is who has the authority to involve law enforcement and when that notification should take place. These are difficult decisions because law enforcement involvement often changes the nature of an investigation and increases the likelihood of public attention.

Communication Plan

As mentioned previously, there are five key stakeholders as part of the incident response. Clearly defined communication protocols and procedures help better prepare the entire team in managing an incident from start to finish and keep everyone apprised of the status.

Key Topic

There are three key things you can do to facilitate and engage with the appropriate members:

Step 1. **Identify the right people.** Find or nominate key individuals within the stakeholder groups. They do not need to be security experts, but they need to be aware of the incident response team's existence. You should make them aware of their duties, which could be to act as a support point for any incident activity.

Step 2. **Set up regular security cadence meetings with all stakeholders.** People forget things, but you can minimize this issue with a regular meeting between all the stakeholders. You can use this meeting to drive improvements, review previous incidents, or plan for exercises.

Step 3. **Escalate the incident response.** When your team is engaged with an incident, you should have them set up proactive alerting. They don't need to call everyone every time, but your handlers need to plan ahead. Your incident response team needs to be warming up key contacts so that when they have to notify them, it doesn't come as a surprise.

Key Topic

Disaster Recovery Plan

A ***disaster recovery plan (DRP)*** is a formal document created by organizations that contains detailed instructions on how to respond to unplanned incidents such as natural disasters, power outages, cyber attacks, or other disruptive events. The plan should contain strategies on minimizing the effects of a disaster so that an organization can continue to operate or be able to quickly resume key operations.

Disruptions lead to lost revenue, brand damage, and dissatisfied customers. And, the longer the recovery time, the greater the adverse impact to the business. This means a good disaster recovery plan should enable rapid recovery from disruptions, regardless of the source of the disruption.

A DRP is more focused than a business continuity plan and does not necessarily cover all contingencies for business processes, assets, human resources, and business partners. A successful DR solution addresses all types of operation disruptions, not just the major natural or person-made disasters that make a location unavailable. Disruptions can include power outages; telephone system outages; temporary loss of access to a facility due to fire, water, or floods; facility threats; or a low-impact nondestructive fire or other event. A DRP should be organized by type of disaster, affected systems, and the location. It must contain instructions that can be implemented by anyone; you cannot always count on the people who put the plan together to be available 24/7 during a disaster. Figure 27-4 illustrates the fundamentals of a successful disaster recovery plan.

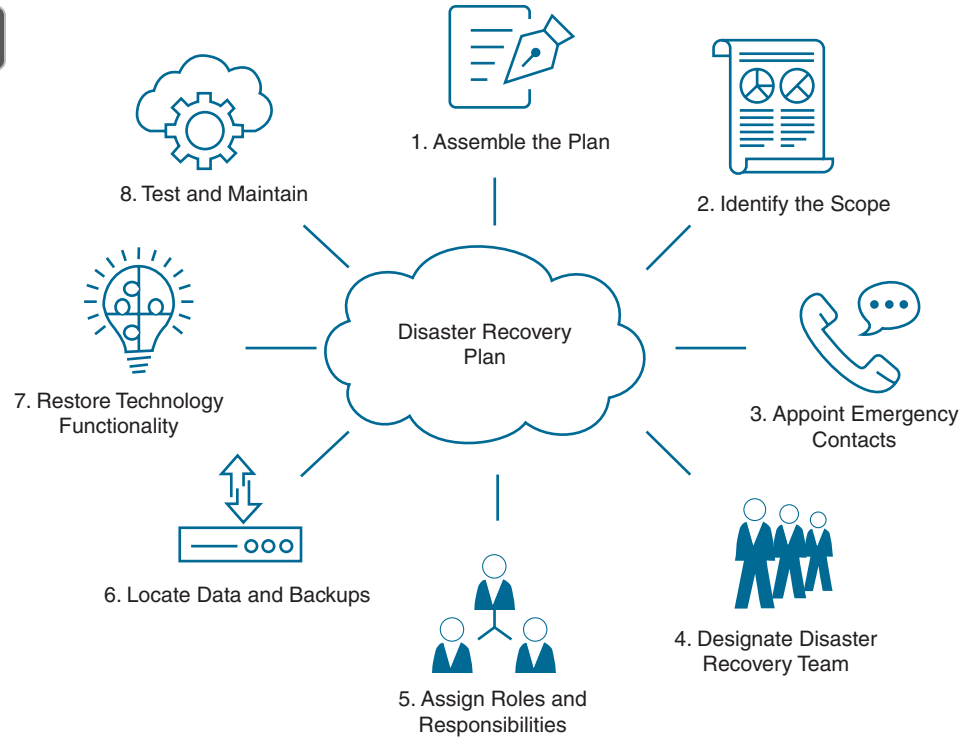


FIGURE 27-4 Disaster Recovery Plan

Disaster recovery planning starts with defining key stakeholders and the data gathering process, a business impact analysis (BIA), risk analysis, recovery strategies, and finally the DR plan. All disaster recovery plans and business impact analyses are cyclical processes that must be maintained and exercised on a regular basis.

Business Continuity Plan

We rarely get advance notice that a disaster is ready to strike. Even with some lead time such as with hurricanes, multiple things can go wrong; every incident is unique and unfolds in unexpected ways. This is where a ***business continuity plan (BCP)*** becomes key to an organization's survival. The appropriate BCP will provide your organization with the best shot at success during a disaster. You need to put a current, tested plan in the hands of all team members and personnel responsible for carrying out any part of that plan. The lack of a plan does not just mean a delay in recovering from an event or incident in your organization; it also could potentially mean the end of the business.

Business continuity (BC) refers to maintaining business functions or quickly resuming them in the event of a disruption, whether caused by a storm, flood, fire, or a malicious attack by cybercriminals. A business continuity plan outlines procedures and instructions an organization must follow in the face of such disasters; it covers business processes, assets, human resources, business partners, and more.

Many people think a DR plan is the same as a BCP, but a DR plan focuses mainly on restoring the IT infrastructure and operations after a crisis. The DR plan is actually one part of a complete business continuity plan, because a BCP looks at the continuity of the entire organization, meaning it needs to be more thorough.

Do you have a way to get Sales, Human Resources, Manufacturing, and Support services functionally up and running so the company can continue to operate and make money after a disaster? For example, if the building that houses your customer service representatives is flattened by a natural disaster like a storm, do you know how those salespeople will be able to handle customer calls? Will you have them work from home temporarily or from an alternate location? The BCP addresses these types of concerns.

A business impact analysis is another part of a BCP. A BIA identifies the impact of a sudden loss of business functions, usually quantified as a direct cost to the business. This analysis can also help you evaluate whether you should outsource non-core activities in your BCP, which can come with its own risks. The BIA essentially helps you look at your entire organization's processes and determine which are most important.



Continuity of Operations Planning (COOP)

Continuity of operations planning (COOP) is a federal initiative to encourage people and departments to plan to address how critical operations will continue under a broad range of circumstances. COOP is important as a good business practice and because the planning fosters recovery and survival in and after emergency situations. A COOP plan addresses emergencies from an all-hazards approach. A continuity of operations plan establishes policy and guidance, ensuring that critical functions continue and that personnel and resources are relocated to an alternate facility in case of emergencies. The plan should develop procedures for

- Alerting, activating, notifying, and deploying employees
- Identifying critical business functions
- Establishing an alternate facility or work-from-home process
- Creating a roster of personnel with authority and knowledge of business operational functions

Creating a continuity of business operations plan is a guided process and a team effort; it will draw on the team's understanding of department operations with emergency management's expertise in preparing for contingencies.

The plan could be activated in response to a wide range of events or situations—from a fire in the building to a natural disaster to the threat or occurrence of a terrorist attack. Any event that makes it impossible for employees to work in their regular facility could result in the activation of the continuity plan.

NOTE Continuity planning is simply the good business practice of ensuring the execution of essential functions and a fundamental duty of public and private entities responsible to their stakeholders.

Key Topic

Incident Response Team

An ***incident response team*** is a group of business and information technology professionals in charge of preparing for and reacting to any type of organizational emergency. Responsibilities of an incident response team include developing a proactive incident response plan, testing and resolving system vulnerabilities, maintaining strong security, and exercising best practices enterprisewide. The team provides support for all incident handling measures, and their expertise typically covers various technical skills, backgrounds, and roles to be prepared for a wide range of unforeseen security incidents.

In incident response, types of emergencies are usually categorized in two ways:

- **Public incidents:** These incidents affect an entire community. They could include natural disasters, terrorist attacks, and widespread epidemics.
- **Corporate/organizational incidents:** These incidents are typically organization-specific and happen on a smaller scale. They could include data breaches, cybersecurity attacks, and physical location threats.

Incident response teams are trained to be prepared for both types and are common in organizations and businesses with valuable intellectual property (IP). As mentioned previously in the chapter, an incident response team could take the following forms:

- **Computer emergency response team (CERT):** This team of professionals is in charge of handling cyber threats and vulnerabilities within an organization. In addition, CERTs tend to release their findings to the public to help other companies and teams strengthen their security infrastructure.

- **Computer security incident response team (CSIRT):** This team of professionals is responsible for preventing and responding to security incidents. A CSIRT may also handle aspects of incident response for other departments, such as dealing with legal issues or communicating with the press.
- **Security operations center (SOC):** Typically, this team includes investigators, threat hunters, and analysts who focus on system security incident response; they often are part of or contribute to the CSIRT investigation.



Retention Policies

An incident response plan must have a data retention policy in preparation for executive review, internal audit, and for possible prosecution. The policy should include specific steps on preserving data and documenting the chain of custody. Without this policy, the cause of the data breach could remain unknown, and a similar breach could occur again in the future. Additionally, an organization can experience numerous legal repercussions for failing to properly preserve data. During a cyber incident, legal counsel should work with the incident response team. The attorney's participation during this process helps establish an attorney-client privilege regarding incident response inquiries, data breaches, or cyber attacks. Note that some tasks, such as when an IRT performs daily operations, do not need or receive attorney-client privilege.

Retention policies for the government agencies fall under NIST SP 800-53, which outlines the requirements contractors and federal agencies need to take to meet the Federal Information Security Management Act (FISMA). It requires data retention for a minimum of three years. SP800-53 also describes how to develop specialized sets of controls tailored for specific types of missions or business functions, technologies, or operation.

NOTE For more information, see https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 27-3 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 27-3 Key Topics for Chapter 27

Key Topic Element	Description	Page Number
Section	Incident Response Plans	761
Figure 27-2	Incident Response Process	765
Section	Tabletop	765
Section	Walkthroughs	766
Section	Simulations	766
Section	MITRE ATT&CK	767
Section	The Diamond Model of Intrusion Analysis	768
Figure 27-3	Diamond Model of Intrusion Analysis	769
List	Description of cyber kill chain attack framework	770
Section	Stakeholder Management	771
List	Facilitating and engaging key stakeholders of the incidence response team as part of a communication plan.	772
Section	Disaster Recovery Plan	772
Figure 27-4	Disaster Recovery Plan	773
Section	Business Continuity Plan	773
Section	Continuity Of Operations Planning (COOP)	774
Section	Incident Response Team	775
Section	Retention Policies	776

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

incident response plan, tabletop exercise, MITRE ATT&CK, Diamond Model of Intrusion Analysis, cyber kill chain, disaster recovery plan (DRP), business continuity plan (BCP), continuity of operations planning (COOP), incident response team

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. Having an incident response plan is imperative; the first step is identifying and having the right people with the right skill sets and experience available and ready to respond. How often should you test and update your plan?
2. Part of your incident response process includes the Eradication phase. During this phase, how long afterward should you increase your monitoring?
3. Incident response simulations are fundamentally about what?
4. Which attack framework emphasizes the relationships and characteristics of four basic components?
5. The cyber kill chain is a series of eight steps that trace stages of a cyber attack. What is step 4?

This page intentionally left blank



This chapter covers the following topics related to Objective 4.3 (Given an incident, utilize appropriate data sources to support an investigation) of the CompTIA Security+ SY0-601 certification exam:

- Vulnerability scan output
- SIEM dashboards
  - Sensor
  - Sensitivity
  - Trends
  - Alerts
  - Correlation
- Log Files
  - Network
  - System
  - Application
  - Security
  - Web
  - DNS
  - Authentication
  - Dump files
  - VoIP and call managers
  - Session Initiation Protocol (SIP) traffic
- syslog/rsyslog/syslog-ng
- journalctl
- NXLog
- Bandwidth monitors
- Metadata
  - Email
  - Mobile
  - Web
  - File
- NetFlow/sFlow
  - NetFlow
  - sFlow
  - IPFIX
- Protocol analyzer output

Using Appropriate Data Sources to Support an Investigation

Long before an incident takes place, an organization should be running baseline network inventory scans and discovery, running vulnerability scans, and collecting and saving logs from systems. All an attacker needs is one vulnerability to get a foothold in your network. That is why at a minimum, you should scan your network at least once a month and patch or remediate identified vulnerabilities immediately. Some compliance requirements require you to scan your network quarterly; however, that is not often enough and should be part of your bigger security commitment. You can't adequately defend your network until you have conducted a security assessment to identify critical assets and know where on your network they are located.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 28-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 28-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Vulnerability Scan Output	1–2
SIEM Dashboards	3–4
Log Files	5–6
syslog/rsyslog/syslog-ng	7
journalctl	8
NXLog	9

Foundation Topics Section	Questions
Bandwidth Monitors	10
Metadata	11
NetFlow/sFlow	12
Protocol Analyzer Output	13

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which is one of the benefits of storing historical vulnerability scans?
    a. Previous scans compared with current scans can provide insight into what has changed.
    b. Historical scans are required for department leads to understand bandwidth issues and plan for upgrades.
    c. Historical scans can be a good data source to feed into your disaster recovery management system.
    d. None of these answers are correct.

2. Which devices on a given network should a vulnerability scanner scan?
    a. Laptops and workstations
    b. Multifunction printers
    c. Routers and switches
    d. All of these answers are correct.

3. When is the best time to implement a SIEM sensor on your corporate networks that support critical assets?
    a. During the incident to capture as much evidence as possible
    b. Before the incident to capture as much evidence as possible
    c. After the incident to capture as much evidence as possible
    d. All of these answers are correct.

4. When you are tuning sensitivity of the SIEM for collecting and alerting on suspect data, what is the risk of an alert being missed because of poor tuning?
   a. Configurations enhance the usability of the SIEM.
   b. The SIEM collects all of the potential incident data required.
   c. A breach can go unnoticed for months.
   d. A breach is alerted and recognized immediately.

5. Which log files that you store in a read-only mode can help with your incident response investigation?
   a. Windows server logs
   b. Linux server logs
   c. Routers and switch log files and logging
   d. All of these answers are correct.

6. Which Session Initiation Protocol is used to establish, maintain, and tear down a call session?
   a. MGCP
   b. RS232
   c. H.323
   d. SIP

7. You need to log data from remote Linux-based applications and devices. Which of the following are your best options? (Select all that apply.)
   a. syslog
   b. rsyslog
   c. syslog-ng
   d. Event Viewer

8. Which **journalctl** command allows you to search through Linux log files and select a specific time window, date, and time?
   a. **journalctl time select**
   b. **sh journalctl file "filename" -select "datatime"**
   c. **journalctl -since YYYY-MM-DD HH:MM:SS**
   d. **journalctl -date -filename**

9. Which of the following is used for centralized logging across various platforms and supports a myriad of different log types and formats?
   a. NXLog
   b. Application log
   c. Security log
   d. System log

10. Which of the following can map out historical trends for capacity planning and quickly identify abnormal usage and top talkers?
   a. Metadata
   b. Dnsenum
   c. Bandwidth monitors
   d. DNSSEC

11. Which type of data does not seem to have any relevance or bearing on an investigation by itself but when combined with additional context is valuable?
   a. Metadata
   b. RFC data
   c. Syslog
   d. SIEM data

12. How does NetFlow create flow sessions from unidirectional sets of packets?
   a. Unidirectional flows are sent to dynamic collectors.
   b. NetFlow uses flows from stateless packets to build an aggregation of flow records.
   c. NetFlow statefully tracks flows or sessions, aggregating packets associated with each flow into flow records.
   d. None of these answers are correct.

13. What tool can you use to capture network traffic and perform analysis of the captured data to identify potential malicious activity or problems with network traffic?
   a. systemd
   b. TTL
   c. Metadata scanner
   d. Protocol analyzer

Foundation Topics

Vulnerability Scan Output

If there were no vulnerabilities within a network or computer system, there would be nothing to exploit, and the network attack surface would be greatly reduced. However, software vulnerabilities always exist and will continue to, because applications are developed by people, and people make mistakes, all of which can allow attackers to compromise networks. Running regular vulnerability scans is just the first step in your defensive posture. A network scan should include all devices with an IP address (workstations, laptops, printers and multifunction printers, IoT devices, routers, switches, hubs, IDS/IPS, servers, wireless networks, and firewalls) and all the software running on them. You should run both authenticated and unauthenticated scans. Each provides insight into vulnerabilities found in services running on your network, open ports on devices that could allow malicious apps to run or communicate on them, and configurations or issues that affect security.

Key Topic

The vulnerability scan reports should be saved for at least 24 months, but certain regulations and industries require longer retention times. Some systems enable you to save these reports and compare month over month to show your cybersecurity posture and improvements. Historical vulnerability scans can also provide significant insight after an incident. By comparing previous scans with the most recent, you can also look for variances in devices and systems that may have been changed. In turn, you should make reimaging these devices a top priority. A change can be anything from additional TCP/UDP ports being shown as open to the detection of unauthorized software, scheduled events, or unrecognized outbound communications.

Figure 28-1 shows basic vulnerability scan results from my internal network 1.1.100.0/24 showing two certificate issues as medium-level threats. Vulnerability reports are perfect for engaging several departments in your organization to get them involved and responsible for the overall posture.

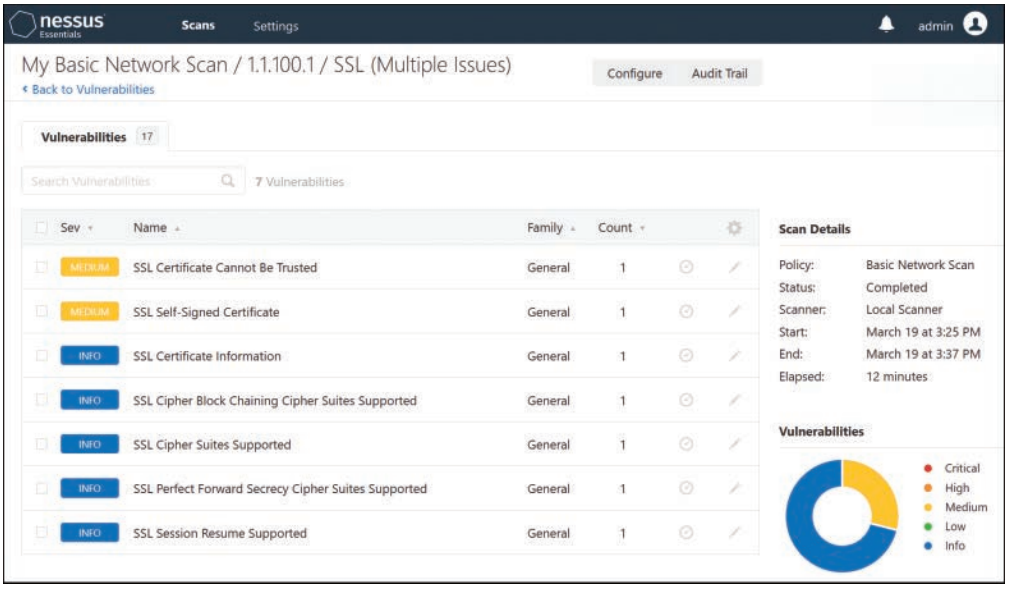


FIGURE 28-1 Vulnerability Scan Report

Key Topic

SIEM Dashboards

Security Information and Event Management (SIEM) software works by collecting logs and event data generated by an organization's application, security devices, and host systems and bringing it together into a single centralized platform, allowing companies to identify threats in real time. SIEM gathers data from antivirus events, firewall logs, and other locations; it then sorts this data into categories. When the SIEM identifies a threat through network security monitoring, it generates an alert and defines a threat level based on predetermined rules.

In Figure 28-2 the SIEM dashboard provides a quick reference to the top threats that require security analyst attention. Alarms by Day Past 30 Days enables you to trend your performance on threats. Top Classification shows types of active attacks in the last nine hours. Top Host (Origin) shows who is attacking and from where, as well as hosts impacted by attacks.



FIGURE 28-2 SIEM Dashboard

TIP SIEMs normalize and aggregate collected data. Log aggregation is the process by which SIEM systems combine similar events to reduce event volume.

Key Topic

Sensors

SIEM sensors collect network traffic and provide advanced threat detection, anomaly detections, log collection, and more. Active network sensors probe and query application scanners and devices on the network for current or existing asset information. Network event sensors provide situational awareness of unauthorized events that take place on the network. Sensors also collect and report installed software information and versions from devices, including patch level. Software and application scanners can be utilized to identify approved and unapproved software installed on devices that may be found through scans. Sensors should be used before and after an incident. Sensor placement around the network allows for greater visibility and can aid in forensic investigation by quickly identifying the breadth of the spread.

Sensitivity

SIEMs have the capability to tune *sensitivity* to what is considered suspect behavior or suspicious files (risk-based prioritization) to help reduce or increase the amount of data/matches and therefore the number of notifications. During an investigation, the more data, the better. On SIEMs and sensors that store lots of data, adjusting the sensitivity to what is considered suspect becomes much easier and can help pinpoint suspect zero. Because of the sensitivity of some company and government data, they are more likely to run the SIEM scans and sensors locally, with little to no external access. More modern hybrid SIEMs with artificial intelligence and machine learning (AI/ML) capabilities can offload some of the more advanced capabilities to these cloud-based services. In this case, a secure tunnel can be established, and only suspect files can be interrogated offsite.

Trends

At the most basic level, SIEMs should enable alerting of suspect traffic, user activity, and files, as well as correlation of data, including multiple sources. They also should allow input of third-party threat intelligence, automatic uses cases, machine learning, and multiple statistical models to be applied to users and network traffic. You should expect to see trend data and more deepening of automatic expert analysis, more automated incident response capabilities, and convergence of network traffic analysis and endpoint detection and response.

NOTE A SIEM dashboard contains multiple views that allow you to visualize and monitor patterns and trends. The intent of a SIEM dashboard is to give you one pane of glass to view your network's or system's status and condition.

Alerts

Alerts provide a purpose for those who monitor SIEMs. When all of the information is flowing into your SIEM, you can track important statistics, identify exceptions, and automate functions like alerting and sending notifications via email, text, or phone, if certain thresholds are met, such as a user logging in with a certain ID or someone accessing certain company assets. Alerting is one of the most important features of a SIEM, and you should take care in how you configure and review alert capabilities.

Correlation

Data *correlation* allows you to take data and logs from disparate systems, such as Windows server events, firewall logs, VPN connections, and RAS devices, and bring

them all together. Because this data is all formatted differently, the SIEM data correlation function allows you to tie these logs together to see exactly what took place during that event.

Log Files

Key Topic

Log files from every system in your network are important. Most devices are capable of sending syslog data of some sort. The log data you collect from your systems and devices might seem pretty mundane; however, these logs could contain the precise evidence needed to investigate and successfully prosecute a crime. For log data to stand up in court as admissible evidence, you must put into place and operate with *chain of custody* in mind and take special precautions on how you collect, handle, and store the data. Many regulations such as PCI DSS, HIPAA, and SOX require the use of logs and log management. When set up properly and with the appropriate due care, logs can provide an immutable fingerprint of system and user activity. In many cases, the logs tell a story as to what really happened in an incident. They can tell what systems were involved, how the systems and people behaved, what information was accessed, who accessed it, and precisely when these activities took place. Log files should be saved to a write once read many (WORM) device like DVD write once (DVD-R) media.

Figure 28-3 shows syslog messages sent from network devices and servers. Notice that the mail server is running postfix/smpd, letting you know this is a Linux host and that it rejected a connection from a host because it could not retrieve its entity information. There are several free visual syslog servers including one from SolarWinds.

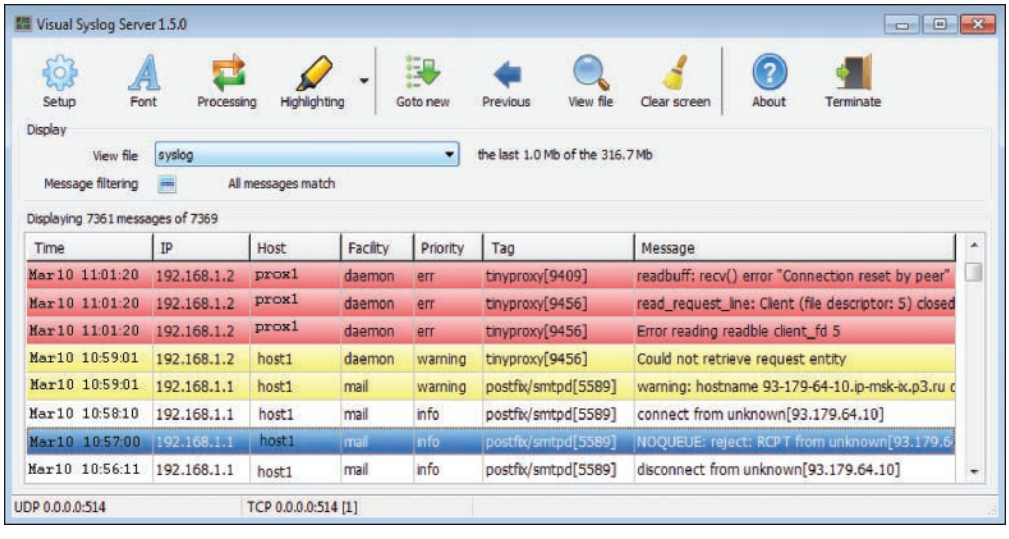


FIGURE 28-3 Visual Syslog Messages

Network

Key Topic

Logging and logfiles for networking devices should have detailed date and time stamps down to the second. You should configure these devices with a common secure Network Time Protocol (NTP) time provider and configure logging to be sent to a separate secure syslog server with limited access. At certain times, logs and logfiles may be available only on the device itself. In those cases, you should ensure access is controlled and logged. These log files may be the key to solving an intrusion or other event.

Every network device—from firewalls to routers to switches to wireless access points—logs slightly differently and at different levels. You should determine what types of logs from these systems you require. Logging everything is not only taxing for the devices but also inundates your logging server. You should take the time to learn the appropriate syntax to set up logging properly for the specific platform. Transit network devices are of particular interest because they are located between your network and your partners, vendors, or some offsite locations. Transit network devices contain syslog data, debugs, and messages that provide additional insight and context during a security incident. This insight aids in determining the validity and extent of an attack.

Within the context of a security incident, investigators can use syslog messages to understand communication relationships, timing, and, in some cases, the attackers' motives and/or tools. The events become part of the larger picture when used in conjunction with other forms of network monitoring that may already be in place. Depending on the device, the logs have a different format and layout, and they provide different information. For example, you can configure logging on a Cisco ASA through the ASDM software for managing ASA firewalls. To do so, you enable Facility 23, timestamp logging, and ensure you have debug-trace logging disabled.

Figure 28-4 shows a real-time log viewer from the ASA ASDM management software. As you can see, the logs are from the ASA firewall, showing that the firewall intercepted a DNS reply for a host on the outside interface and providing both IP addresses. This information allows you to dig deeper in the Syslog Details window.

RAW syslog messages from devices look similar to the syslog details shown in Figure 28-4, although using Linux command-line utilities or GUI-based tools is usually best. One of the reasons logging to SIEMS is so helpful is that it can understand the context of what is being received and has the ability to correlate, deduplicate, and provide a clear picture of the event.

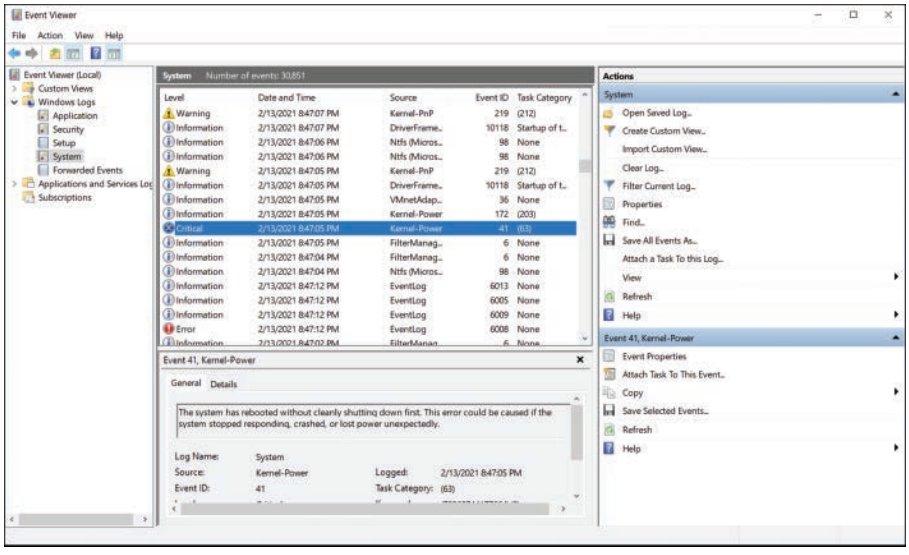


FIGURE 28-4 Real-Time Log Viewer

Key Topic

System

System logging is all about the capabilities of the specific system you are trying to obtain logs from. Each system that has logs also refers to specific system messages that contain log records for the operating system events. Many of them show how system processes and drivers were loaded, and errors during loading or any failures since. Windows provides a few areas to obtain logs, such as the Event Viewer. You can use it to collect .evt and .evtx files from other Windows-based machines. You can even sort by event ID and perform filtering and exporting. The Windows Admin Center can provide events from Event Viewer as well and allows you to export them.

Linux hosts log files are mostly contained in the /var/log directory. You should make it a regular process to collect these logs and store them as you would with other devices. Application, system, and security logs should be stored off-host when possible. Monitoring system logs provides proactive monitoring and helps you build a complete picture if or when an incident occurs: knowing which systems were accessed and when, what files were accessed, which applications were installed, and whether the integrity of the system was compromised and how.

Figure 28-5 shows the Windows Event Viewer where you can clearly see the event ID, level of logging, and type of message. Here, four types of messages are displayed: Warning, Critical, Information, and Error. Information messages are typically those from the operating system letting you know that it performed some action like allowing a process to start. Errors are generated by failures of processes to

start. Critical messages indicate some serious failure; in this case, the Event Viewer assumes the system did not shut down properly. Warning messages are normal; Windows handles all of these events and recovers without any user intervention.

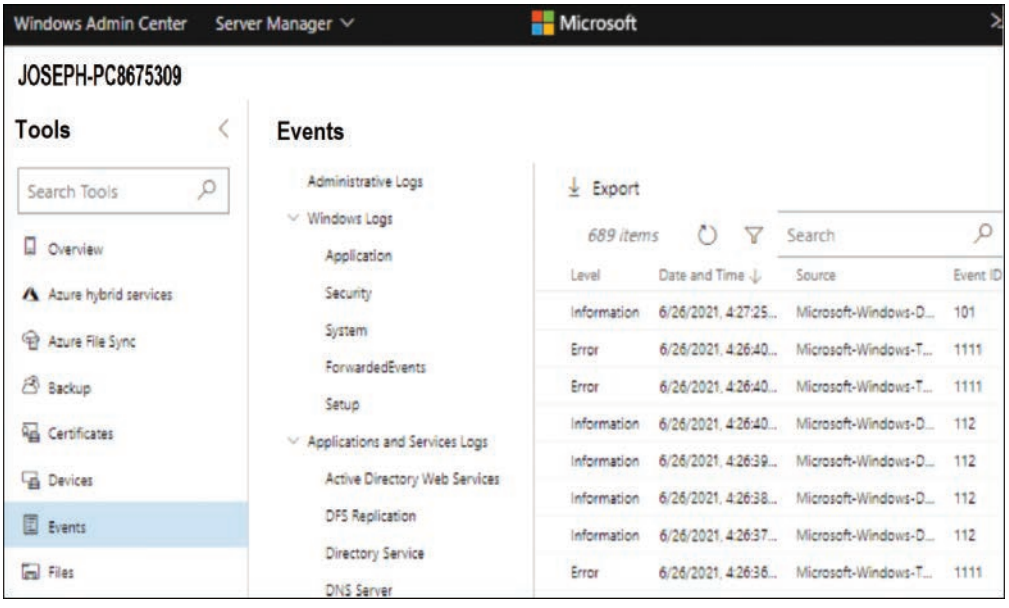


FIGURE 28-5 Event Viewer

The system log records error messages, warnings, and other information generated from the Windows operating system. It provides a wealth of information on what is going on with the system. System log records can also help you determine if, when, or where a user attempted to log in and failed, what user ID was used, and what IP address it came from.

Key Topic

Application

Logging hosts end to end from services, events, or systems to applications gives you near end-to-end visibility. Application logging and the process for implementation, management, and reporting need to be configured and reviewed before, during, and after an incident. You should parse, partition, and analyze all logs and data generated by the application. Logging for critical process information about user, system, and web application behavior can help incident responders build a better understanding of the breadth and depth of an attack. Application logs provide information related to applications run on the local system, how they are performing, and any attempts to utilize the application. Depending on the level of logging, the logs can determine whether users

exposed an unknown flaw in the application programming or whether they attempted to escalate privilege or modify data they were not intended to have access to.

Key Topic

Security

Both successful and failed login events are relevant to logging and help build the investigation's scope. Security event–specific logs from systems, devices, and applications allow you to focus on main event types. They help guide the investigation and allow the team to quickly get the organization to the breadth of the breach, attribution, and recovery. Computer forensics processes use log file data in finding electronic evidence for criminal investigations. Companies should follow best practices to help ensure log data and log management practices properly support forensic investigations.

The security log in the Windows Event Viewer provides information related to the success and failure of login attempts and information related to other audited events, as shown in Figure 28-6.

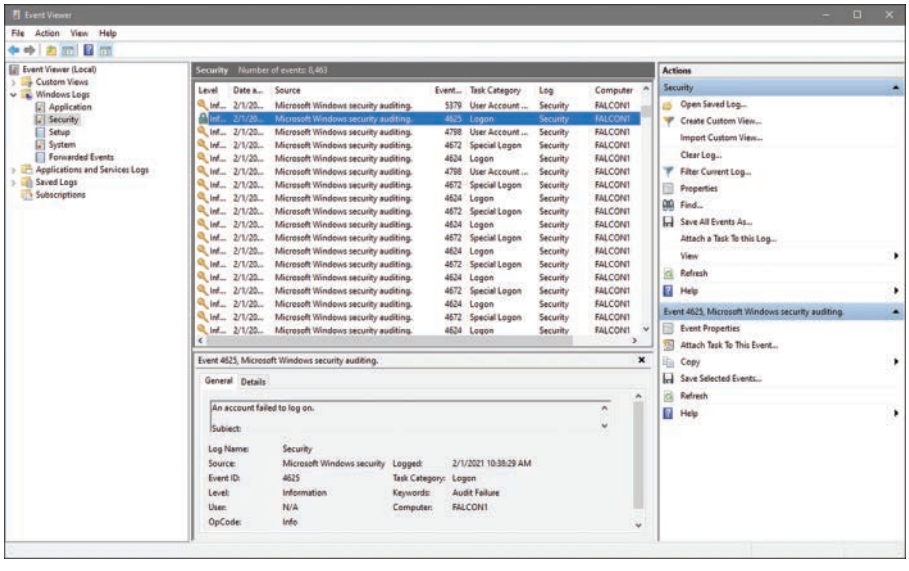


FIGURE 28-6 Security Log Event Viewer

NOTE Windows Event Viewer logs are stored in %SystemRoot%\System32\Winevt\Logs.



Web

Log files provide a precise view of the behavior of the server as well as critical information about when, how, and by whom a server is being accessed. This kind of information can help incident responders and investigators unfold the chain of events that may have led to a particular activity. For example, the most common web servers are IIS, Nginx, and Apache. These servers provide log files such as access.log and error.log. The access.log records all requests for files made by visitors. For example, if a visitor requests www.example.com/main.php, an entry is added in the log file showing what the visitor requested. The log files provide investigators with an IP address, date, and time, as well as requested files. In the absence of that log file, you might have never known that someone discovered and ran a secret or restricted script you have on your website that dumps the database.

Analyzing logs from web servers such as Apache, Nginx, and IIS can provide important insight into the website and web application quality and availability. Web server logs are usually access logs, common error logs, custom logs, and W3C logs. W3C logs are used mainly by web servers to log web-related events, including web logs. The log format for an Apache access log file is as follows:

```
LogFormat "%h %1 %u %t \%r\" %>s %b" common
```

Logs can be an extremely important aspect of your web environment. They provide additional data that's useful for security, debugging purposes, informational purposes, and more.

If a request was made to a website using the log format, the resulting log would look similar to the following.

```
127.0.0.1 - chris [26/Feb/2021:10:34:12 -0700] "GET /sample-image.png
HTTP/2" 200 1479
```

Figure 28-7 shows a web server's log file, where users are requesting specific pages. Based on the GET requests and 404 messages, the user is requesting files that do not exist, and this could be an attacker looking for potential files to exploit or an outdated webpage somewhere making a reference to files that have been since removed.

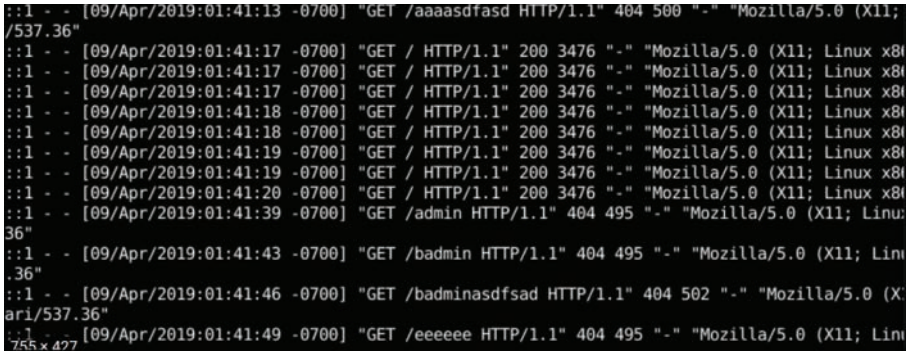


FIGURE 28-7 Web Server Log File

Key Topic

DNS

The Domain Name System (DNS) provides a hierarchy of names for computers and services on the Internet or other networks. Its most noteworthy function is the translation of domain names such as example.com into IP addresses. DNS is required for the Internet to function, operates on a global scale, and is massively distributed. Authoritative DNS servers run a service called BIND that normally accepts messages on UDP port 53. BIND performs the DNS role acting as an authoritative name server for domains and also acts as a recursive resolver in the network. The DNS protocol has two message types: queries and replies. Both use the same format. These messages are used to transfer resource records (RRs). An RR contains a name, a time-to-live (TTL), a class (normally IN), a type, and a value.

There are nearly a dozen different types of logs that are of particular interest. Obtaining and including them in your investigation can help build a full picture. Investigating any transfers or lookups and any queries prior to the attack can help identify a source IP address. DNS servers can and should be configured for file integrity monitoring systemwide. Figure 28-8 demonstrates how to use the **tail** command to view the /named/security.log, showing a DNS client request and DNS A record being approved for entry.

```
root@kilo:/var/log/named# tail /var/log/named/security.log
23-Oct-2021 15:56:54.171 security: debug 3: client 127.0.0.1#37
957: request is not signed
23-Oct-2021 15:56:54.171 security: debug 3: client 127.0.0.1#37
957: recursion available
23-Oct-2021 15:56:54.171 security: debug 3: client 127.0.0.1#37
957: query (cache) 'lamsecurs.info/A/IN' approved
23-Oct-2021 15:56:54.171 security: debug 3: client 127.0.0.1#37
957: request is not signed
23-Oct-2021 15:56:54.172 security: debug 3: client 127.0.0.1#37
957: recursion available
23-Oct-2021 15:56:54.172 security: debug 3: client 127.0.0.1#37
957: query (cache) 'lamsecurs.info/A/IN' approved
root@kilo:/var/log/named#
```

FIGURE 28-8 Security.log

The following are common DNS record types:

- **Address mapping record (A record):** Also known as a DNS host record, stores a hostname and its corresponding IPv4 address. When you request google.com, it provides you with an associated IP address.
- **IP version 6 address record (AAAA record):** Stores a hostname and its corresponding IPv6 address.

The Domain Name System Security Extensions (DNSSEC) protocol was developed to strengthen DNS through the use of digital signatures and public key cryptography. It protects against attacks by providing a validation path for records.

The DNS enumeration tool called dnsenum enumerates DNS by finding DNS servers and DNS records such as mail exchange servers, domain name servers, and the address records for a domain. You can use nslookup and dig to troubleshoot DNS servers and resolve name resolution issues by querying a DNS server to check whether the correct information is in the zone database. For Windows, you use nslookup to check which DNS servers are being used, and you use dig, which stands for Domain Information Groper, on Mac and Linux-based systems.

Several log aggregators for DNS Bind, such as Men & Mice Logeater, can provide query statistics that are formatted to make it easy to understand loads and issues on DNS servers. These tools are also capable of grouping error messages, even with DNSSEC.



Authentication

One of the primary logs that should be used in any incident response or forensic investigation is the authentication log. Well-written applications also log authentication approved and failure events. By reviewing authentication failures, you can see password guessing and crackers at work. If you were to search even deeper and correlate which users logged in and authenticated and when they did so, then compare with valid user locations and the systems they were actually on, and filter on those suspect differences, you could build a picture of which systems the attackers gained access to and what activity was performed. To view authentication pass/fail attempts on Linux, you navigate to the /var/log/ folder and locate the auth.log file, where you can use **cat** or **more** on the file. Receiving an Open response means that someone logged on. In Figure 28-9, the wrong password was entered, which shows up as an authentication failure. The figure also shows that the session opened for user root, which indicates a successful login attempt.

```
Mar 21 14:03:15 amx5 sudo: pam_unix(sudo:auth): authentication failure; logname= uid=100
0 euid=0 tty=/dev/pts/1 ruser=adminx rhost=  user=adminx
Mar 21 14:03:17 amx5 sudo: pam_unix(sudo:auth): Couldn't open /etc/securetty: No such fi
le or directory
Mar 21 14:03:41 amx5 sudo: pam_unix(sudo:auth): Couldn't open /etc/securetty: No such fi
le or directory
Mar 21 14:03:41 amx5 sudo:   adminx : TTY=pts/1 ; PWD=/root ; USER=root ; COMMAND=/usr/b
in/su root
Mar 21 14:03:41 amx5 sudo: pam_unix(sudo:session): session opened for user root by (uid=
0)
Mar 21 14:03:41 amx5 su: (to root) root on pts/1
Mar 21 14:03:41 amx5 su: pam_unix(su:session): session opened for user root by (uid=0)
```

FIGURE 28-9 Authentication Log

Key Topic

Dump Files

Everyone has seen the blue screen of death. Once it is initiated on a Windows system, it creates a dump file that is sent to %systemroot%\Minidump. These dump files contain information that was in memory on the system right before it died. This information includes applications, processes, and activities, including commands run before the crash. All of these indicators can help you build a picture of what an attacker was doing on the system before it crashed. In many cases, you can see programs the attacker was running, and what is even more helpful is that the image can contain evidence before the attacker had an opportunity to clean up. Tracking the specific programs and commands used by a would-be attacker helps create an attacker profile.

To access the dump file, first you need to enable the setting to write debugging information. As shown in Figure 28-10, you can select Small Memory Dump, which makes it easier to read with standard tools. After this setting is enabled, if you have a crash, you can easily find the file in %systemroot%\Minidump.

As you can see in Figure 28-11, on your Windows host, if your system crashed or you initiated a dump file, you see a MEMORY.DMP file under c:/windows/Minidump/. Using a tool called Bluescreen View, you can dive deeper into the reason for the failure. In this case, it's obvious that the problem was a driver because IRQ_NOT_LESS indicates a card or driver conflict.

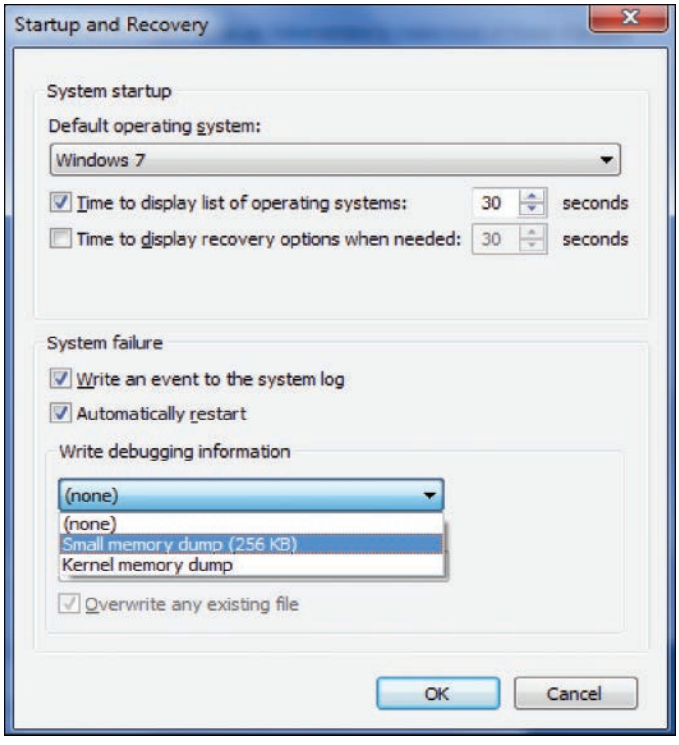


FIGURE 28-10 Enabling the Setting to Write Debugging Information

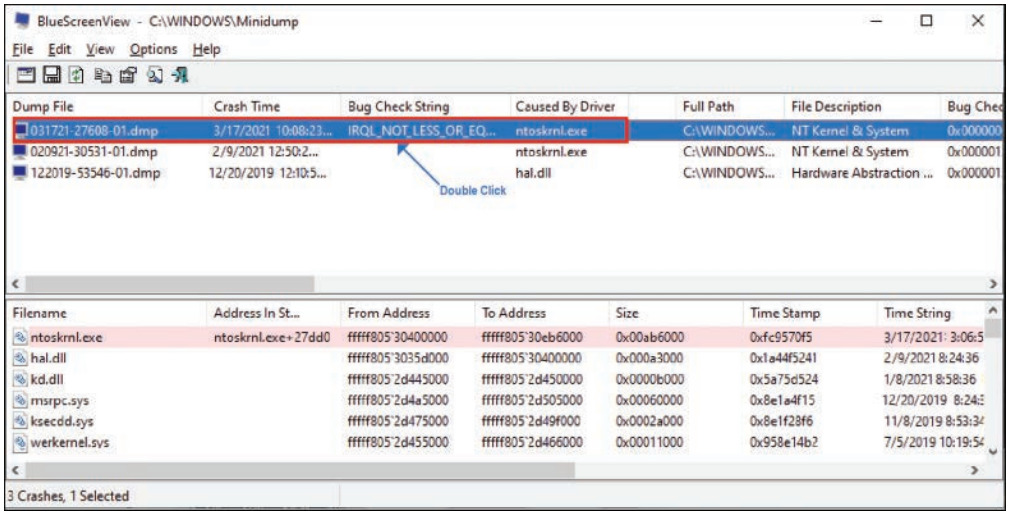


FIGURE 28-11 Viewing Dump File Information

Key Topic

VoIP and Call Managers

The popularity of Voice over Internet Protocol (VoIP) is increasing as the cost savings and ease of use are realized by a wide range of large and small company users. VoIP technology also is an attractive platform to criminals. The reason is that call managers and VoIP systems are global telephony services, so it is difficult to verify the user's location and identification. The security of placing such calls is also appealing to criminals; they can make calls and use systems as pivot points to attack other systems.

As shown in Figure 28-12, the logging and logs from these systems can provide the context of attacker calls made, including source and destination of IP addresses, as well as any systems connected to. Several tools on Call Manager allow you to view calls and troubleshoot. The Real-Time Monitoring Tool (RTMT) provides access to the syslog viewer. Ideally, you can configure RTMT to send all logs to a remote syslog server, which ensures the attacker isn't able to modify logs that may indicate the intrusion.

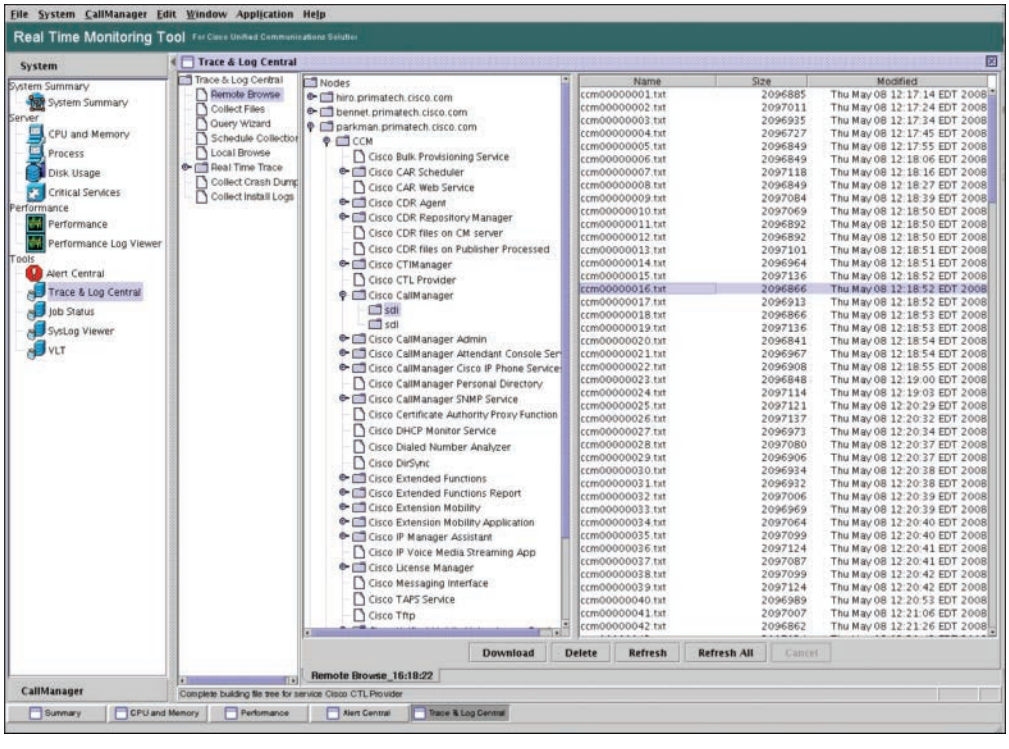


FIGURE 28-12 Call Manager Logs

If you have Call Detail Records (CDR) enabled, you can go to **CDR > Search > By User/Phone Number/Sip URL** and select Next. Input the extension in the Phone Number/SIP URL field and click the Add button. Then select the time and date range and click Find, export in CSV, and save to an XLS; then you can filter your data.

Key Topic

Session Initiation Protocol Traffic

Session Initiation Protocol (SIP) is a signaling protocol used to establish, maintain, and tear down a call when terminated. SIP allows the calling parties called user agents (UAs) to locate one another using the network, which allows user agents to be registered and invite other UAs to join in an Internet multimedia call called a session. SIP devices are typically routers connected to the telco or to the Internet; these devices should be configured to send detailed time-stamped logs to a syslog server. Routers can be used as pivot devices to gain additional access to the network; therefore, hardening the routers is important. Combining logs from these devices with other logs collected from call systems allows the investigator to build a picture of what attackers were doing, who they were contacting, and what systems they compromised. Depending on your SIP provider and what system you are using to process SIP calls, your logging can be different. Most SIP providers include tools that provide reporting from call diagrams to failed attempts and even alerting. Figure 28-13 shows some sample SIP logs.

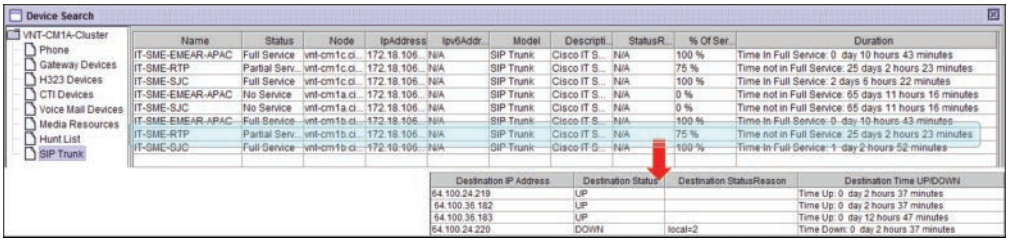


FIGURE 28-13 SIP Logs

Key Topic

syslog/rsyslog/syslog-ng

The system logging service ***syslog*** was created in the 1980s to provide logging for systems. The ***syslog-ng*** is an advancement to syslog that added many enhanced features, such as content-based filtering, direct logging to a database, using TCP for transport, and TLS for encryption. All of these improvements are important for protecting information being sent from systems and devices. The rocket-fast syslog server, or ***rsyslog***, was created in 2004, extending the syslog protocol with buffered operation and Reliable Event Logging Protocol (RELP) support.

Logging to an external system is a very important aspect of maintaining chain of custody; it also ensures the logs of the attacked device are not tampered with because they are stored off-system. Although syslog is a decades-old standard for message logging, it is still extensively used. It is available on most network devices (such as routers, switches, and firewalls), as well as printers and Linux-based systems. Over a network, a syslog server listens for and then logs data messages coming from the syslog client.

Some Debian-based Linux systems such as Ubuntu store global system activity and startup messages in /var/log/syslog.

rsyslog and syslog-ng are similar because they build on syslog capabilities by adding support for advanced filtering, configuration, and output. syslog-ng is considered next-gen, a free and open-source implementation of the syslog protocol for Linux systems. It extends the original syslogd model with content-based filtering, rich filtering capabilities, and flexible configuration options. It also adds important features to syslog, such as using TCP for transport. syslog-ng started from scratch, with a different configuration format, whereas rsyslog was originally a fork of syslogd, supporting and extending its syntax. syslog-ng provides the capability for many third-party tools to interface and manipulate data, including Grafana, splunk, logzilla, and others. They provide a web interface and allow you to perform robust searches using correlation and advanced queries, as well as provide many other open-source options for viewing and manipulating log files.

Figure 28-14 shows the command-line help interface for syslog-ng.

```
└─# syslog-ng --help                                                    1 ×
Usage:
  syslog-ng [OPTION?] syslog-ng

Help Options:
  -h, --help                                        Show help options
  --help-all                                        Show all help options
  --help-process                                    Process options
  --help-log                                        Log options

Application Options:
  -V, --version                                     Display version number (sys
log-ng 3 (3.28.1))
  --module-path=<path>                              Set the list of colon separ
ated directories to search for modules, default=${exec_prefix}/lib/syslog-ng/3.28
  --module-registry                                 Display module information
  -S, --seed                                        Does nothing, the need to s
eed the random generator is autodetected
```

FIGURE 28-14 syslog-ng Command-Line Help Interface

journalctl

Key Topic

journalctl is a Linux command-line tool used for viewing logs that are collected by systemd. These logs are collected in a central repository for easy display, review, and reporting. The log records are well indexed and structured in a way that enables you to easily analyze and manipulate the large log records. A number of filtering options enable users to sift through large sets of log entries to find specific data fields, messages, alerts, and errors that are specific to certain applications, users, services, and applications. To run journalctl on a Debian distribution, you need to install **systemd**, **apt-get install systemd**. Figure 28-15 illustrates the journalctl reviewing log files from this system showing log start date and the version of Linux. This is what a raw log format looks like.

```
-- Journal begins at Mon 2021-03-01 16:05:26 CST, ends at Sun 2021-03-21 16:21:30 CDT. >
Mar 01 16:05:26 amx5 kernel: Linux version 5.9.0-kali1-amd64 (devel@kali.org) (gcc-10 (>
Mar 01 16:05:26 amx5 kernel: Command line: BOOT_IMAGE=/boot/vmlinuz-5.9.0-kali1-amd64 r>
Mar 01 16:05:26 amx5 kernel: Disabled fast string operations
Mar 01 16:05:26 amx5 kernel: x86/fpu: Supporting XSAVE feature 0x001: 'x87 floating poi>
Mar 01 16:05:26 amx5 kernel: x86/fpu: Supporting XSAVE feature 0x002: 'SSE registers'
Mar 01 16:05:26 amx5 kernel: x86/fpu: Supporting XSAVE feature 0x004: 'AVX registers'
Mar 01 16:05:26 amx5 kernel: x86/fpu: Supporting XSAVE feature 0x008: 'MPX bounds regis>
Mar 01 16:05:26 amx5 kernel: x86/fpu: Supporting XSAVE feature 0x010: 'MPX CSR'
Mar 01 16:05:26 amx5 kernel: x86/fpu: xstate_offset[2]:  576, xstate_sizes[2]:  256
Mar 01 16:05:26 amx5 kernel: x86/fpu: xstate_offset[3]:  832, xstate_sizes[3]:   64
Mar 01 16:05:26 amx5 kernel: x86/fpu: xstate_offset[4]:  896, xstate_sizes[4]:   64
Mar 01 16:05:26 amx5 kernel: x86/fpu: Enabled xstate features 0x1f, context size is 960>
Mar 01 16:05:26 amx5 kernel: NX (Execute Disable) protection: active
Mar 01 16:05:26 amx5 kernel: SMBIOS 2.7 present.
Mar 01 16:05:26 amx5 kernel: DMI: VMware, Inc. VMware Virtual Platform/440BX Desktop Re>
Mar 01 16:05:26 amx5 kernel: vmware: hypercall mode: 0x00
Mar 01 16:05:26 amx5 kernel: Hypervisor detected: VMware
Mar 01 16:05:26 amx5 kernel: vmware: TSC_freq read from hypervisor : 3096.000 MHz
Mar 01 16:05:26 amx5 kernel: BIOS-e820: [mem 0x00000000bfeff000-0x00000000bfefffff] ACP>
```

FIGURE 28-15 Journalctl Raw Log Format

To access the man (full manual) page, you enter **man systemd.journal-fields**.

To access logs for a specific time window, you type **journalctl -since '2019-03-11 15:05:00'**. The time format is YYYY-MM-DD HH:MM:SS.

As you can see, journalctl is a valuable tool you can use to collect logs and investigate a Linux system as well as other syslog records/files to sort through mountains of data to help you find the proverbial needle in a haystack. This tool enables you to query and display logs from journald, which is the logging service for Linux-based systems that stores log data in binary format. With journald, you cannot log to remote locations; however, there are several mechanisms to forward logs to syslog when necessary.

NOTE The man page for using Journalctl is located at www.freedesktop.org/software/systemd/man/journalctl.html.

Key Topic

NXLog

NXLog is used for centralized logging across various platforms and supports a plethora of different log types and formats. NXLog is available in two versions: the Community Edition and the Enterprise Edition. Enterprise Edition features are platform support for Linux, Windows, Android, AIX, Solaris, and macOS. NXLog can accept event logs from TCP, UDP, files, databases, and various other sources in different formats such as syslog or Windows event logs. NXLog can process high volumes of event logs from many different sources. Log processing includes rewriting, correlating, alerting, filtering, pattern matching, log file rotation buffering, and prioritized processing. NXLogs can store or forward event logs in a number of supported formats.

To use NXLog on Windows, you must make sure you use the appropriate x86 version and use an MMC snap-in (such as gpedit.msc).

On Debian Linux type operating systems like Ubuntu, you download the appropriate .deb file and install. To start NXLog on Linux, you enter **service nxlog start**.

Many investigators use NXLog to uncover supportable evidence, enabling you to build a near complete picture of the status of the host. Because it can prioritize certain critical log messages, it ensures those attacker messages make it to NXLog (central log collection). A common requirement is to detect conditions when there are no log messages coming from a source. This usually indicates a problem, such as network connectivity issues, a server that is down, an unresponsive application, or an attacker cleaning up. The absence of logs can also be a good reason to investigate.

The solution to this problem is to use statistical counters and scheduled checks. The NXLog input module can update a statistical counter configured to calculate events per hour. In the same input module, a schedule block checks the value of the statistical counter periodically. When the event rate drops below a certain limit, an action can be executed, such as sending out an alert message. In Figure 28-16, the Windows entry for nxlog.log shows the startup of NXLog and the IP address and port being used.

nxlog.log

```
2021-1-16 16:53:10 INFO nxlog-5.2.6573 started
2021-1-16 16:53:10 INFO connecting to 192.168.40.43
2021-1-16 16:53:12 INFO successfully connected to 192.168.40.43:1514
2021-1-16 16:53:12 INFO successfully connected to agent manager at 192.168.40.43:4041 in SSL mode
```

FIGURE 28-16 The nxlog.log File Provides Insight into the Startup

Key Topic

Bandwidth Monitors

A ***bandwidth monitor*** tracks bandwidth use over all areas of the network, including devices, applications, servers, WAN, and Internet links, and that information will assist you in keeping an eye on inbound and outbound bandwidth within your network and help you identify which hosts are using the most bandwidth. One benefit of deploying bandwidth monitors is that they map out historical trends for capacity planning. With bandwidth monitors, you can quickly identify abnormal bandwidth usage, top talkers, and unique communications, all useful in finding infected systems that may be exfiltrating data or scanning the network looking to spread to other hosts. Bandwidth monitors provide critical information before, during, and after investigations. Incident responders can use this baseline information to determine when the attacked host started to overcommunicate outbound or to spread internally. The historical information is key to determining "normal" for the attacked network and host, what is normal communication for that specific host, and what is normal for the network in general. There are several developers of bandwidth monitors, and some devices have built-in bandwidth logging and monitoring.

In Figure 28-17, the SolarWinds monitor shows several problem devices, including which stations are using the most bandwidth, where they are going, and how much traffic (both in megabytes and in percentage). Several baseline items show how normal traffic should look as compared to what is currently happening.

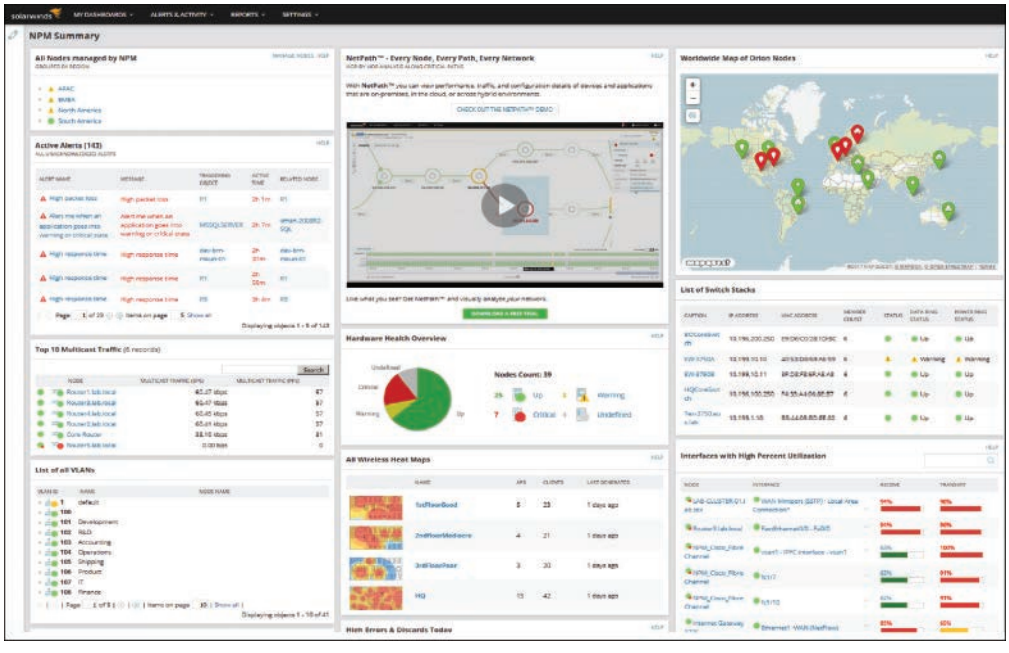


FIGURE 28-17 Bandwidth Monitoring

Metadata

Metadata is created from every activity you perform, whether it's on a personal computer or online, every email, web search, and social or public application and interaction. Metadata is defined as "data that provides information about other data." There are many distinct types of metadata. On its own, it may have little or no relevance to the investigation. However, when placed with other types of data and sources, the result could bring you steps closer in your investigation. Metadata security is the practice and policies designed to protect an organization from security risks posed by unauthorized access to the organization's metadata. Even metadata from Microsoft Word documents contains the names of authors and modifiers, dates of creation, and changes and file size. Metadata risk is the potential for disclosure of private information, usually unknowingly, because this metadata is hidden from plain view and users may be unaware of it. This information may be something those outside your company shouldn't have access to.

There are several types of metadata you should be aware of. The first is descriptive metadata, which refers to elements like titles, dates, and keywords. For instance, when you download a video file, it contains descriptive metadata describing what the video is about as well as the name and date. Descriptive metadata is used with books, providing book titles, author names, dates, pages, and key data locations.

In addition to descriptive metadata, structural metadata provides information concerning a specific object or resource. Digital media (film on DVD, for example) is a perfect example. Each section has a certain length of film running time. In a broader sense, structural metadata records information on how a particular object or resource might be sorted.

Preservation metadata provides information that strengthens the entire process of making a digital file (object). This includes vital details required for a system to communicate or interact with a specific file and upholds the integrity of a digital file or object. A common pattern involved is the Preservation Metadata Implementation Strategies (PREMIS) mode that brings forth common factors to preservation and maintenance, including actions taken on a digital file or the rights attached to it.

Use metadata is data that is sorted each time a user accesses and uses a specific piece of digital data. Use metadata is gathered in a clear and direct attempt to make potentially helpful predictions about a user's future behavior. This type of metadata can be used to understand fluctuations in data that have no pattern when there is really a pattern beneath the data.

There are even more types of metadata, including provenance. This means that it's most relevant when something changes or is duplicated frequently, such as in the digital realm where there are frequent changes. Administrative metadata informs

users what types of instructions, rules, and restrictions are placed on a file. This information helps administrators limit access to files based on user qualifications.

Cell phones have a ton of metadata, from the pictures you take that provide GPS coordinates, to the date and time of the photo, the language, camera type, flash setting, and more. From the coordinates of the tweet or Instagram image you just sent, to metadata from emails, and public and private social media applications you use, to airline tickets and rental car details, your phone is literally a treasure trove of information.

Figure 28-18 shows a Microsoft Word file. By right-clicking and looking at the properties under Details, you can see the title, authors, subject, tags, and comments. Notice the Remove Properties and Personal Information link, which allows you to remove certain personal information.

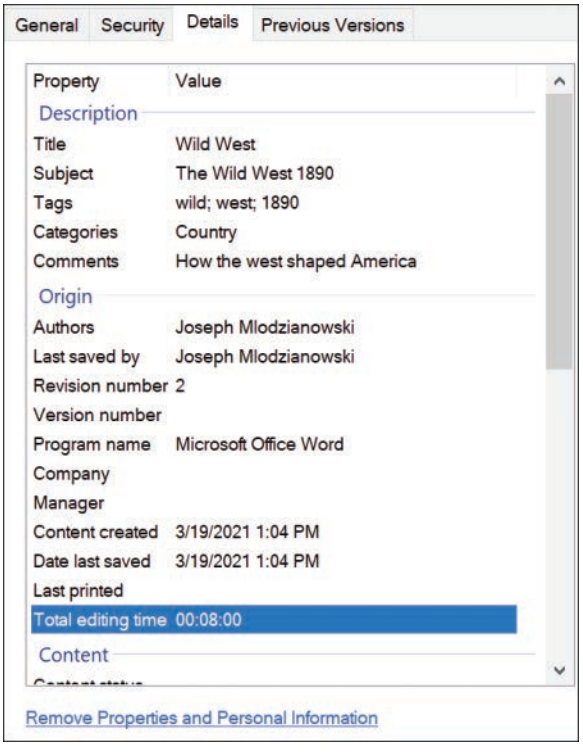
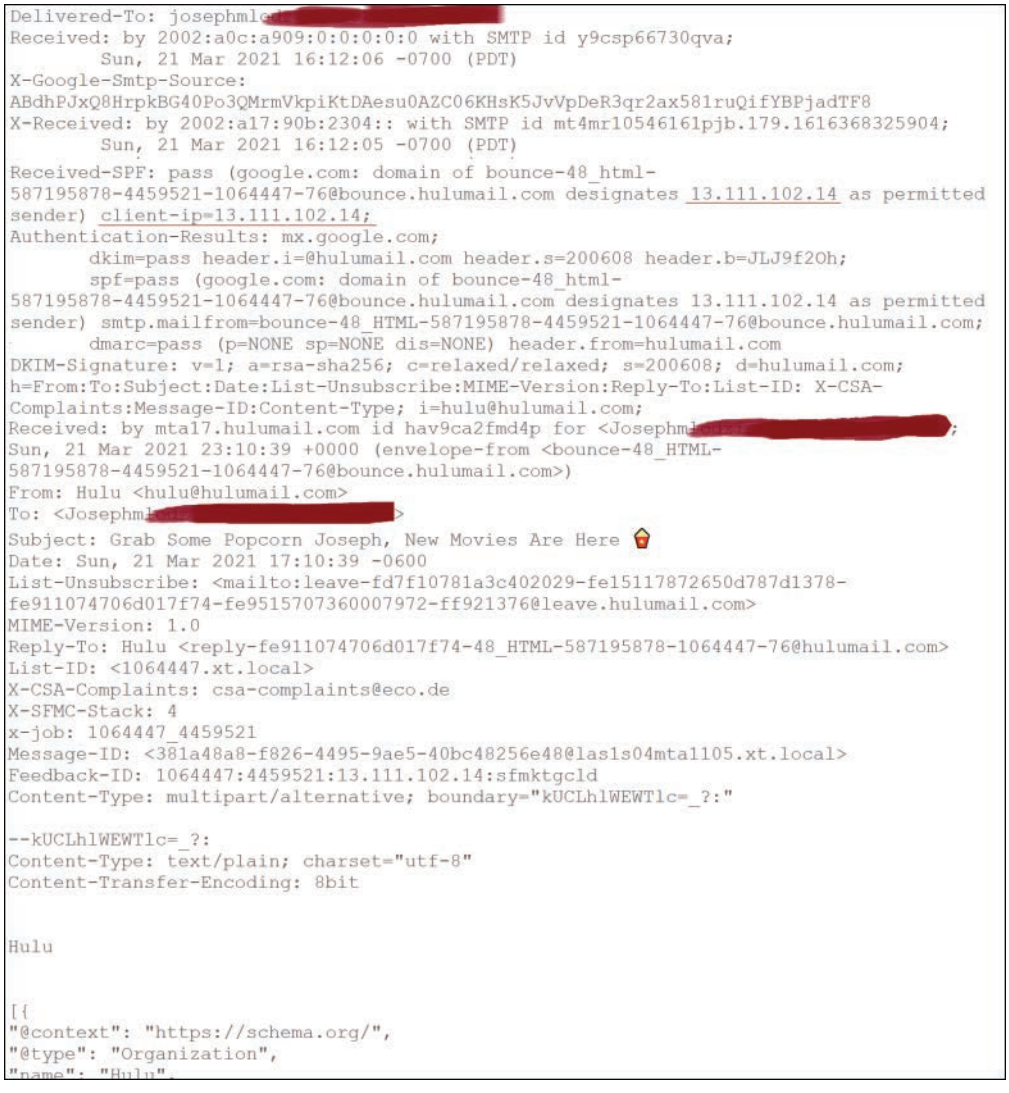


FIGURE 28-18 Metadata in a Microsoft Word File

Figure 28-19 shows an email header's metadata about an email message received from Hulu at hulumail.com. As you can see, this message provides email addresses for the server (Google server), how the email was handed off, and the path it took to get to my Gmail account (gmail email mx.google.com) server. So, grab some popcorn and look through the email headers to see what else you can find of interest.



```
Delivered-To: josephmlg[redacted]
Received: by 2002:a0c:a909:0:0:0:0:0 with SMTP id y9csp66730qva;
        Sun, 21 Mar 2021 16:12:06 -0700 (PDT)
X-Google-Smtp-Source:
ABdhPJxQ8HrpkBG40Po3QMrmVkpiKtDAesu0AZC06KHsK5JvVpDeR3qr2ax581ruQifYBPjadTF8
X-Received: by 2002:a17:90b:2304:: with SMTP id mt4mr10546161pjb.179.1616368325904;
        Sun, 21 Mar 2021 16:12:05 -0700 (PDT)
Received-SPF: pass (google.com: domain of bounce-48_html-
587195878-4459521-1064447-76@bounce.hulumail.com designates 13.111.102.14 as permitted
sender) client-ip=13.111.102.14;
Authentication-Results: mx.google.com;
       dkim=pass header.i=@hulumail.com header.s=200608 header.b=JLJ9f2Oh;
       spf=pass (google.com: domain of bounce-48_html-
587195878-4459521-1064447-76@bounce.hulumail.com designates 13.111.102.14 as permitted
sender) smtp.mailfrom=bounce-48_HTML-587195878-4459521-1064447-76@bounce.hulumail.com;
       dmarc=pass (p=NONE sp=NONE dis=NONE) header.from=hulumail.com
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/relaxed; s=200608; d=hulumail.com;
h=From:To:Subject:Date:List-Unsubscribe:MIME-Version:Reply-To:List-ID: X-CSA-
Complaints:Message-ID:Content-Type; i=hulu@hulumail.com;
Received: by mta17.hulumail.com id hav9ca2fmd4p for <Josephml[redacted]>;
Sun, 21 Mar 2021 23:10:39 +0000 (envelope-from <bounce-48_HTML-
587195878-4459521-1064447-76@bounce.hulumail.com>)
From: Hulu <hulu@hulumail.com>
To: <Josephml[redacted]>
Subject: Grab Some Popcorn Joseph, New Movies Are Here
Date: Sun, 21 Mar 2021 17:10:39 -0600
List-Unsubscribe: <mailto:leave-fd7f10781a3c402029-fe15117872650d787d1378-
fe911074706d017f74-fe9515707360007972-ff921376@leave.hulumail.com>
MIME-Version: 1.0
Reply-To: Hulu <reply-fe911074706d017f74-48_HTML-587195878-1064447-76@hulumail.com>
List-ID: <1064447.xt.local>
X-CSA-Complaints: csa-complaints@eco.de
X-SFMC-Stack: 4
x-job: 1064447_4459521
Message-ID: <381a48a8-f826-4495-9ae5-40bc48256e48@las1s04mta1105.xt.local>
Feedback-ID: 1064447:4459521:13.111.102.14:sfmktgcld
Content-Type: multipart/alternative; boundary="kUCLhlWEWTlc=_?:"

--kUCLhlWEWTlc=_?:
Content-Type: text/plain; charset="utf-8"
Content-Transfer-Encoding: 8bit


Hulu


[{
"@context": "https://schema.org/",
"@type": "Organization",
"name": "Hulu",
```

FIGURE 28-19 Full Headers from a Gmail Email

NOTE Metadata has proven itself to be very useful in regards to investigations, and practically everything digital contains metadata including files, images, and email.

Email

There is a ton of metadata in emails, specifically in the headers. Email is quite trivial to spoof and/or trick users into clicking on. Once attackers have a foothold in a network, they can turn any host into a spammer. There is a lot you can learn from these emails and those that might be used in spear phishing executives of the company.

The header is a section of code that contains information about where the email came from and how the message reached its destination. A header contains the originator's email address and/or the computer that the perpetrator/sender was using, the browser used, and the IP address, which is sometimes conveniently identified as the Originating IP. With this information, you can trace the email to the Internet service provider (ISP) with the date and time of the email using the sender's computer's IP address.

There are two types of email headers. There are partial headers that you see when you normally open your email. The partial headers are the most important to your daily tasks. Such headers are the From Address, To Address, Subject, Date and Time, Reply-To Address, CC, and BCC. And then there are full headers; these are simply more technical information than you normally see when you check your email. They are required to complete a full forensic investigation. Figure 28-19 shows only a small part of full headers; full headers allow you to forensically follow the patch from the sender to the receiver.

Email is responsible for over 85 percent of successful company intrusions, so you should ensure you have detailed logging enabled.

Mobile

Mobile devices are essentially people's lives now. We pay our bills, chat with friends, interact on social media platforms, perform online banking, and much more. Using a mobile device to obtain metadata can provide a significant advantage to the incident response team. If the phone received a text message, IM, email, or other type of communication and it was suspect zero, detailed mobile forensic applications can export and provide these files in an investigative form. Beyond that, you can use the metadata from the device to see if any precipitating activity led to the attack/breach.

Web

Web pages are a consistent source of metadata. While not usually visible, it's in the keywords and phrases that describe the contents of the page. Most of this metadata is picked up and used by search engines to index a site. Search engine optimization (SEO) practices are those that take advantage of meta tags, such as keywords, phrases, page titles, and more. During an incident, investigators should review all

web page metadata content to see if it was removed or modified. There have been incidents where attackers change this data to replace it with vulgar, adult, or malware links or words. Then it gets picked up by search engines and scored low or hostile, which makes the company website unavailable to browsers and search engines because of the ratings.

File

Certain applications add attributes to the files they create or edit. Accessing this metadata can be an important step in finding who created the files, what application was used, the host system operating system, and much more. For any file you have on your computer, simply right-click on it, select Properties, and then select Details. What you see then is all metadata, and most of it is not readily visible to the user.

Key Topic NetFlow/sFlow

NetFlow and sFlow are technologies that can aid in gathering additional information about network, user, and application traffic that can be used to build a better understanding of network, user, and application flow.

NetFlow

NetFlow is a session flow protocol that collects and analyzes network traffic data that can be used to help you understand which applications, users, and protocols might be consuming the most network bandwidth or if a denial-of-service (DoS) activity is taking place and who the actors are. A "flow" is a unidirectional set of packets sharing common session attributes, such as source and destination, IP, TCP/UDP ports, and type of service. NetFlow statefully tracks flows or sessions, aggregating packets associated with each flow into flow records, which are bidirectional flows.

The export of NetFlow records depends on timers. It may take up to 30 minutes to export flows. If NetFlow is deployed on a compromised network, the flow sensors and flow tools can aid in the investigation by providing details of the communication between the attacker and the investigated device, as well as the tactics, techniques, and procedures (TTPs) of the attacker. This is another tool in the incident response team's toolchest that provides valuable information.

Because NetFlow is a protocol developed by Cisco, it is deployed on Cisco routers and switches that are configured to send flow data ultimately to a NetFlow collector. As of this writing, the latest version of NetFlow is NetFlow v9. NetFlow collects packets via router interfaces, which are then provided to a NetFlow collector. A NetFlow application can then analyze the data, which is useful for incident

investigations for intrusion analysis. Figure 28-20 shows NetFlow traffic details. As you can see, it is full of rich and useful information, including application usage, most accessed locations of your network devices, the type of traffic, web browsing, database access, and DNS—all very useful for detecting an attack or network issue.

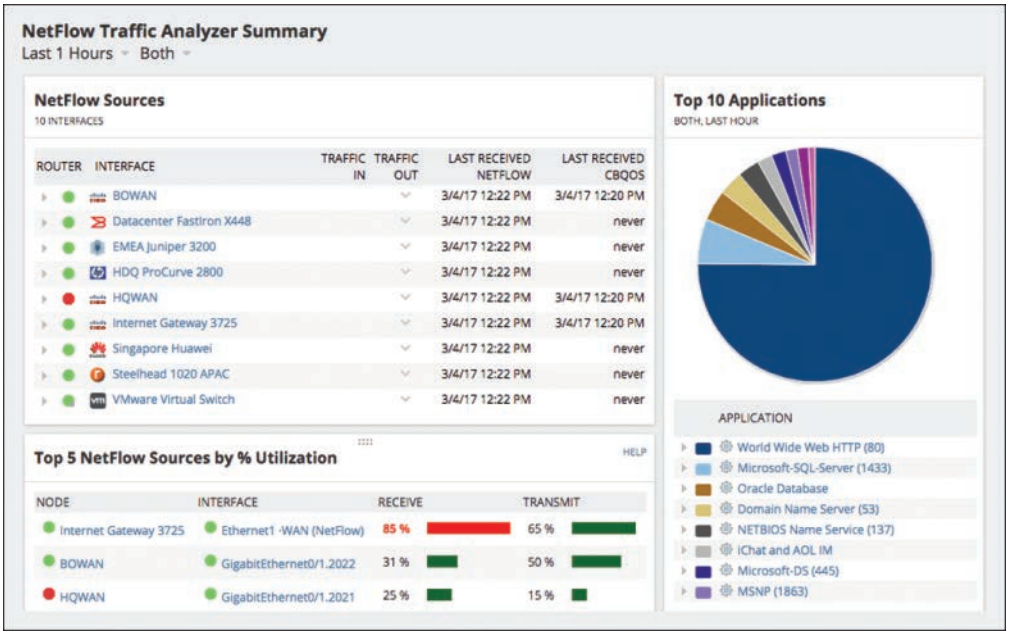


FIGURE 28-20 NetFlow Report

sFlow

Often network engineers look at ***sFlow*** as the successor to NetFlow; however, there are places to deploy specific tools for specific outcomes. sFlow can monitor Layer 2 to Layer 7 headers. sFlow stands for Sampled Flow, where packets are being sampled instead of flows. The goal of packet sampling and filtering is to forward only certain packets. sFlow datagrams are continuously sent across the network in real time. In terms of visibility, sFlow is good at massive DoS attack detection because the sampled network patterns are sent on the fly to the sFlow collector. sFlows are only an approximation of the real traffic because sampled packets do not reflect all network traffic. As a result, sFlow lacks the accuracy that is provided by NetFlow.

sFlow can provide the incident response team with much more insight into network Layers 2–7, which includes the packet header, MPLS labels, VLANs, MAC address, and all of the Layer 3–4 information from NetFlow, providing additional context to the investigation and visibility at more layers. sFlow is used for network monitoring.

However, sFlow samples packets and can provide analysis to identify unauthorized network activity and quickly identify attacks, as well as investigate DDoS attacks.

In Figure 28-21, an sFlow trend shows that traffic slows at each interval, allowing you to quickly identify abnormal traffic and investigate those items.

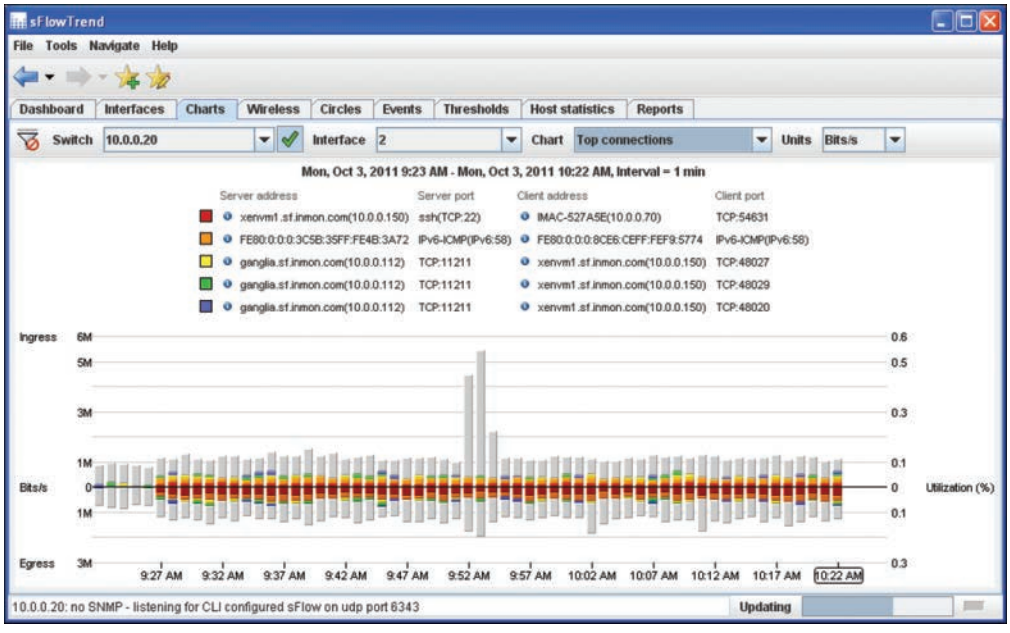


FIGURE 28-21 sFlow Trend Report

Key Topic

IPFIX

Internet Protocol Flow Information Export (IPFIX) is a standard for exporting the information about network flows from devices. Defined in RFC 5982, IPFIX is derived from NetFlow v9. IPFIX has a metering process that generates flow records collecting data packets at an observation point, filters them, and aggregates information about these packets. Flow records are encapsulated by Layer 4 protocols (SCTP, UDP, or TCP) to a collector. IPFIX can be used to export any traffic information from L2 to L7 to a flow collector. IPFIX tracks IP actions across the network. When it collects data packets from across the network, the exporter organizes and sends the compiled information to a collector. In IPFIX, exporters can transport data to multiple collectors, which is known as a many-to-many relationship. Exporters send information sets via IPFIX messages.

IPFIX is a flexible protocol that supports variable-length fields. One major advantage IPFIX has over NetFlow is in the collection process of the data packets, where users can organize and analyze the data using IPFIX. Users are able to customize their requests, so the system completes only specified tasks, such as organizing information or doing basic data analysis. IPFIX is also capable of integrating more information into its exporting process.

When you look at individual host flow ratios, TCP flags, host reputation, and DNS and compare trend data, this flow data can be very effective at detecting malware. NetFlow and IPFIX can be leveraged to study network behaviors over time. Any communication considered abnormal can trigger events that increase indexes, which can trigger alarms and notification. From there, investigators can take over, and with the information collected from IPFIX, they can get a picture of what happened and how.

Figure 28-22 shows the robust data available. IPFIX enables variable-length strings, allows a vendor to be specified, and is an enabler technology that allows vendors to export any performance details they can provide.

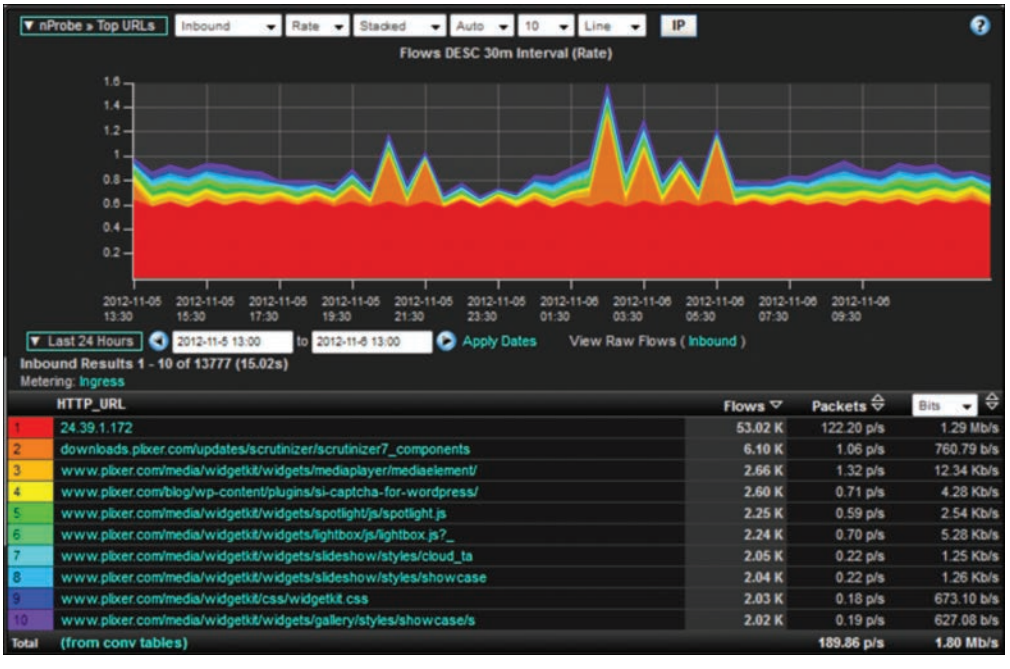


FIGURE 28-22 IPFIX Report

NOTE IPFIX is based on a NetFlow implementation and serves as an industry standard for the export of flow information from all kinds of network devices. It is an IEFT protocol created out of the need to facilitate measurement, billing, and accounting services. You can find more information about IPFIX at https://tools.ietf.org/html/rfc7011.

Protocol Analyzer Output

Key Topic

Protocol analyzers primarily allow network engineers and security teams to capture network traffic and perform analysis of the captured data to identify potential malicious activity or problems with network traffic. The network traffic data can be observed in real time for troubleshooting purposes, monitored by an alerting tool such as a SIEM to identify active network threats, and/or retained to perform forensic analysis. Generally, there are two categories of protocol analyzers.

Ad hoc protocol analysis tools like Wireshark are used to troubleshoot or analyze a specific issue. Although these tools are best focused on analysis for a specific protocol or host, they are not suitable for monitoring an entire network for an extended period.

A number of enterprise analysis tools are more well suited for monitoring your entire corporate infrastructure around the clock. Well-known monitoring solutions such as Paessler PRTG, SolarWinds, Manage Engine, and Atera provide holistic monitoring of protocol traffic, database health, application availability, server uptime, and other collection points. The output from these analyzers can be fed into certain SIEMs and other security tools to not only alert of an attack but also provide proactive notifications using algorithms that can observe the steps taken before an attack.

A protocol analyzer is a tool for monitoring and analyzing capture signals as they traverse communication channels, such as satellite, computer bus, leased line, the Internet, and many other mediums. Figure 28-23 shows a 1480A physical device that is connected via USB to a computer and software that decodes the communication; this one is by International Test Instruments Corporation. This protocol analyzer can capture and decode even low-level bus communication protocols.

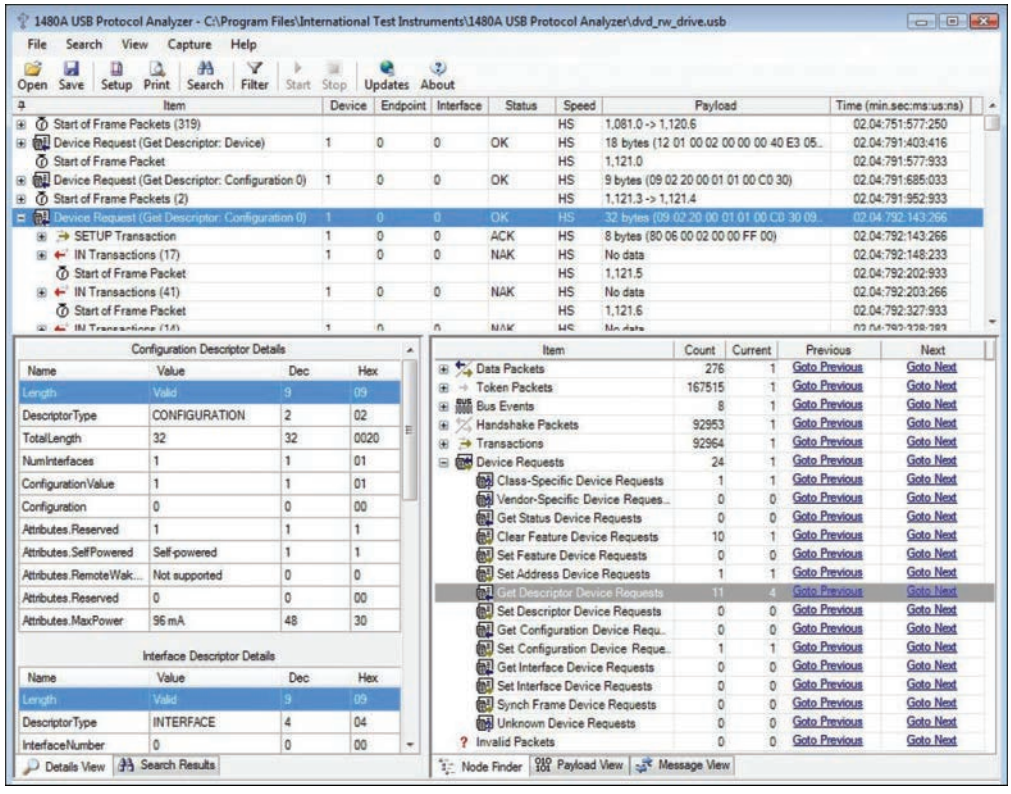


FIGURE 28-23 Protocol Analyzer Report

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 28-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 28-2 Key Topics for Chapter 28

Key Topic Element	Description	Page Number
Paragraph	Vulnerability scan output	785
Section	SIEM Dashboards	786
Section	Sensors	787
Section	Sensitivity	788
Section	Trends	788
Section	Alerts	788
Section	Correlation	788
Paragraph	Log files	789
Section	Network	790
Section	System	791
Section	Application	792
Section	Security	793
Section	Web	794
Section	DNS	795
Section	Authentication	796
Section	Dump Files	797
Section	VoIP and Call Managers	799
Section	Session Initiation Protocol Traffic	800
Paragraph	syslog/rsyslog/syslog-ng	800
Paragraph	journalctl	802
Paragraph	NXLog	803
Section	Bandwidth Monitors	804
Section	Metadata	805
Section	NetFlow/sFlow	809
Section	IPFIX	811
Paragraph	Protocol analyzers	813

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

Security Information and Event Management (SIEM), log files, syslog, syslog-ng, rsyslog, journalctl, NXLog, bandwidth monitor, metadata, NetFlow, sFlow, Internet Protocol Flow Information Export (IPFIX), protocol analyzers

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is an indicator of a host on your network being compromised?
2. How does a SIEM use data correlation to help with discovering what took place during an incident?
3. Application logging can help an investigator by building a picture of what _____ looks like.
4. There are nearly a dozen logs available from DNS. They include which two message types?
5. What platform would cybercriminals use to make nearly anonymous calls?

This page intentionally left blank



This chapter covers the following topics related to Objective 4.4 (Given an incident, apply mitigation techniques or controls to secure an environment) of the CompTIA Security+ SY0-601 certification exam:

- Reconfigure endpoint security solutions
  - Application approved list
  - Application block list/deny list
  - Quarantine
- Configuration changes
  - Firewall rules
  - MDM
  - DLP
  - Content filter/URL filter
  - Update or revoke certificates
- Isolation
- Containment
- Segmentation
- SOAR
  - Runbooks
  - Playbooks

Applying Mitigation Techniques or Controls to Secure an Environment

After a security incident has been stabilized, you, as security administrator, should examine lessons learned to prevent reoccurrences of similar incidents. This includes patching server vulnerabilities, restricting access, training employees, or potentially rolling out technologies to better monitor, alert, and respond to threats. Fixing security flaws or vulnerabilities found during post-incident activities should be a priority, but you should apply mitigation techniques immediately. Finally, you should update your security incident response plan to reflect all these preventive measures and ensure security systems are tuned to avoid missing similar events.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 29-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 29-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Reconfigure Endpoint Security Solutions	1–3
Configuration Changes	4–5
Isolation	6
Containment	7
Segmentation	8
SOAR	9

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Lists are used in every manner of computer and application access. When you allow an application in your system, what type of list do you use?
   a. Approved list
   b. Block list
   c. Allow list
   d. Deny list

2. A block list or deny list is a basic access control capability that denies specific types of what from executing?
   a. LDAP directories
   b. Applications
   c. API interfaces
   d. Scheduled events

3. What is the main function of the quarantine capability?
   a. To safely submit random files to the cloud
   b. To safely copy suspect files
   c. To safely delete suspect files
   d. To safely store suspect files

4. Which attack vector(s) needs to be considered for mobile platforms?
   a. Hardware
   b. Firmware
   c. Mobile OS
   d. All of these answers are correct.

5. Data loss prevention is a type of which of the following?
   a. Corporate protection of physical assets
   b. Assistance with storage of data
   c. Loss prevention process
   d. End-to-end goal

6. Which of the following is a reason you should isolate a host from the network and other hosts?
   a. Virus, Trojan, or other attack
   b. Operating system defect
   c. Downgrading the OS system
   d. Downloading updates

7. Which of the following are not part of the containment phase?
   a. Change passwords.
   b. Disable network connectivity.
   c. Turn off the display.
   d. Place the infected host in a sandbox.

8. Network segmentation is an approach that divides a network into multiple networks. Which of the following is a benefit of a segmented network?
   a. Smaller broadcast domain
   b. Fewer hosts per segment
   c. Ability to apply granular policies
   d. All of these answers are correct.

9. Which system alerts for suspected phishing emails, endpoint attacks, failed user logins, malware, and other threat information come from a variety of detection sources, such as SIEMs, systems, switches, and logging services?
   a. SOAR
   b. Runbook
   c. Playbook
   d. Content filter

Foundation Topics

Reconfigure Endpoint Security Solutions

Endpoints are among the biggest targets for cybercriminals because these devices are especially vulnerable to various attacks. Criminals often attempt to use endpoint devices as entry points to access corporate networks and steal data, leverage existing software vulnerabilities, or hold information hostage. The most effective practices to detect, respond, and protect your endpoints from these threats include a layered defense-in-depth approach.

Application Approved Lists

Key Topic

Antivirus, firewall, and endpoint security management systems generally have application approved lists, sometimes called application whitelists. If you are currently using a block list, you might consider also implementing an application approved list, one that is centrally managed, which gives you more granular control. ***Application approved lists*** are where you spend the time up front building the allow list of approved applications. Then with your central management or endpoint security solution, you roll out the allow list enterprisewide. After an incident, you should spend the time building or updating an application approved list and then enable and enforce it. It will keep users from installing unauthorized applications that could lead to breaches. Figure 29-1 shows a typical approved list in Windows. Each application that is loaded or installed shows up on the list, and you can also import applications. Selecting the box to the left of the application allows you to deny or approve the application.

Application Block List/Deny List

Key Topic

An ***application block list*** or ***deny list*** is a basic access control mechanism that denies specific applications or code on the list from being installed or run. After an incident, part of the mitigation strategies should be to review all rogue and unauthorized applications that were installed by the attacker. You can start with these reports to help build your block list/deny list; there are also public block lists you can use to add to your own list. The problem with block lists is the upkeep and maintenance. It's a simple process for attackers to use renamed or rehashed applications to bypass these lists.

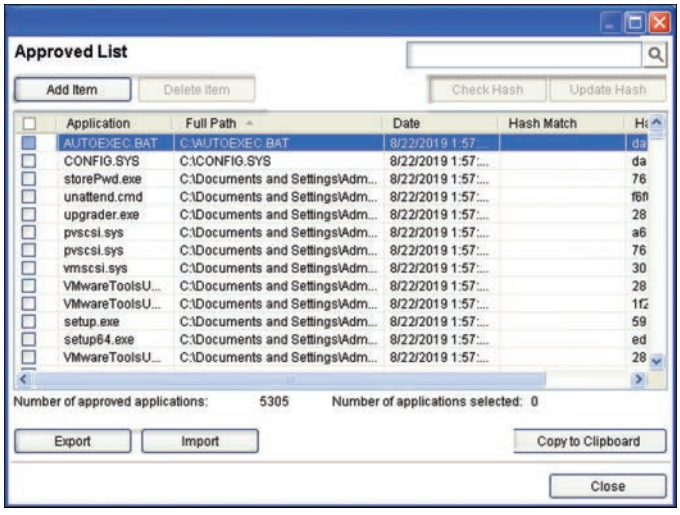


FIGURE 29-1 Windows Application Approved List

TIP A block list or deny list provides a list of all applications that the organization deems undesirable or banned and prevents those applications from being installed. An allow list approach uses a list of approved applications. If the application is not on the approved list of software, the application installation is denied or restricted.

Key Topic

Quarantine

The ***quarantine*** function is to safely store reported objects such as malware, infected files, or potentially unwanted applications. If Microsoft Defender Antivirus is configured to detect and remediate threats on your device, it quarantines suspicious files. By default, the Microsoft Defender Antivirus virus quarantine storage is located under the following path: C:\ProgramData\Microsoft\Windows Defender\Quarantine. Defender creates a threat log during an investigation. This should be one of the places investigators check first for evidence and history of suspect files.

Reverse-engineering these quarantined files can provide insight into the attacker who created, modified, or used the files. After the investigation has wrapped up, the security team should review all files that were quarantined and all those that were missed but eventually were found; reviewing them provides an opportunity to increase the efficacy of the suspect files and increase quarantine. Each antivirus program quarantines its files in different locations. You need to research the antivirus application to locate the files but should be careful when interacting with files in this location.

Figure 29-2 shows the Microsoft Defender Antivirus interface for virus and threat protection, currently running a scan. With managed settings, you can configure multiple items and view logs.



FIGURE 29-2 Microsoft Defender Antivirus Interface for Virus and Threat Protection

Key Topic

Configuration Changes

A flaw in your organization's security settings, like failing to secure your infrastructure, could leave your entire network and every device connected to it vulnerable to an attack. Configuration vulnerabilities are easily preventable provided secure best practices are followed. Applying mitigation techniques after an incident requires a thorough assessment of all network devices.



Firewall Rules

The firewall was introduced back in 1988, and currently, there are many brands and types of firewalls on the market. These firewalls differ not only in cost but also their functionalities. The firewall at the perimeter of a network is the first line of defense against external attacks. To mitigate the attacks, the firewall divides the network into two zones:

- **Trusted Zone:** Authorized users are in a private network or access a private network.
- **Least-Trusted Zone:** Users from the Internet are trying to access the private network.

The job of the firewall is to either permit or deny the traffic based on the *firewall rules*. Firewalls filter packets and mitigate attacks; by utilizing best practices and consistent audits before and after attacks, you can reduce exposure and risk. First, you should utilize the National Institute of Standards and Technology (NIST) SP 800-41 guidelines on firewalls. Firewall policies are addressed in Chapter 4 of that document: Chapter 4.1 describes policies based on IP addresses and protocols; Chapter 4.2 addresses policies based on applications; and Chapter 4.3 addresses policies based on user identity. Adherence to the rule of zero trust and least privilege should always be the goal.

TIP The NIST SP800-41 guidelines for firewalls are located at https://csrc.nist.rip/library/alt-SP800-41-rev1.pdf.



MDM

Mobile device management (MDM) provides functionality and control over enterprise devices by enabling the enrollment of enterprise devices for management functions, such as provisioning devices, tracking inventory, changing configurations, updating, managing applications, and enforcing policies. Unified endpoint management (UEM) expands on MDM by providing unified capabilities and also including content and threat management, identity and access management, and application containerization.

Many organizations administer devices and applications using MDM products/services. MDM primarily deals with segregating corporate data, securing emails, securing corporate documents on devices, enforcing corporate policies, and integrating and managing mobile devices, including laptops and handhelds of various categories. MDM implementations may be either on-premises or cloud-based.

Following are some of the core functions of MDM:

- Ensuring that diverse user equipment is configured to a consistent standard/supported set of applications, functions, or corporate policies
- Updating equipment, applications, functions, or policies in a scalable manner
- Ensuring that users use applications in a consistent and supportable manner
- Ensuring that equipment performs consistently
- Monitoring and tracking equipment (such as location, status, ownership, activity)
- Being able to efficiently diagnose and troubleshoot equipment remotely

MDM functionality can include over-the-air distribution of applications, data, and configuration settings for all types of mobile devices. Most recently, laptops and desktops have been added to the list of systems supported as mobile device management becomes more about basic device management and less about the mobile platform itself. MDM tools are leveraged for both company-owned and employee-owned devices across the enterprise or mobile devices owned by consumers. Consumer demand for bring your own device (BYOD) is now requiring a greater effort for MDM and increased security for both the devices and the enterprise they connect to, especially because employers and employees have different expectations concerning the types of restrictions that should be applied to mobile devices.

Mobile devices are essentially general-purpose computing platforms. They are not restricted to performing one operation and can instead be used in many different fields including medical, industrial, and entertainment. Commercially available mobile devices lack a unified set of features. Each new feature or capability has the potential to introduce new threats to security and privacy. It is important to establish a baseline understanding of the changes that are common to mobile devices. Mobile platforms have several attack vectors that need to be considered: hardware, firmware, mobile OS, applications, and the device combination. Figure 29-3 highlights the various attack vectors on a mobile device.

TIP NIST SP800-124 Rev. 2, Guidelines for Managing the Security of Mobile Devices in the Enterprise is located at https://csrc.nist.gov/publications/detail/sp/800-124/rev-2/draft.

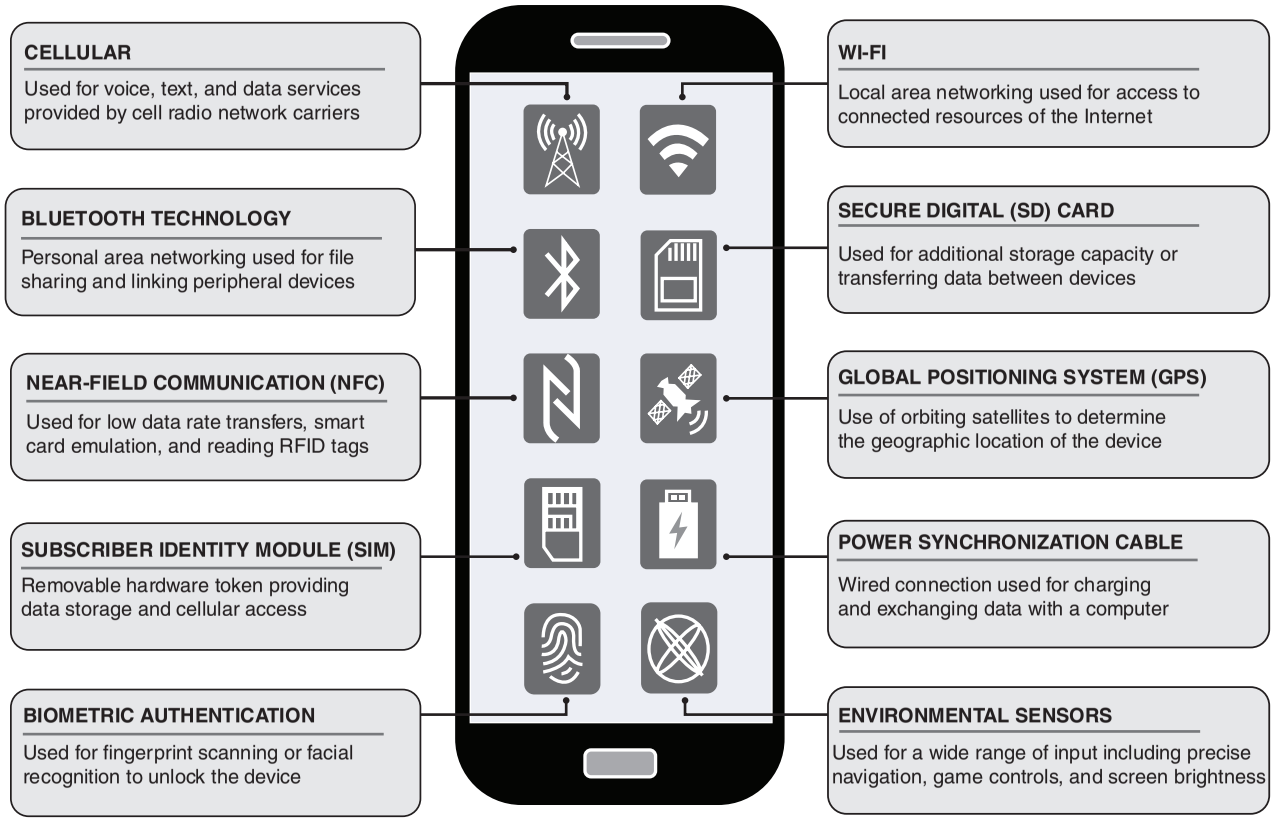


FIGURE 29-3 Mobile Device Attack Surface

Data Loss Prevention

Data loss prevention (DLP) is an end-to-end goal to make sure that users do not send sensitive or critical information outside the corporate network. The term routinely describes software products that help a network administrator control the data that users can view or transfer. Intellectual property, corporate data, and customer data are some of the types of data you would use a DLP system or software to help protect against exfiltration. These systems, if configured correctly, can also alert you to an unauthorized person (attacker) attempting to transfer data that is classified as sensitive. After an incident, if it is determined that the DLP system not only missed the removal of sensitive documents but did not alert you, you will need to review the policies, restrictions, and alerts that are currently configured and review the attack account to determine what was removed and then document, reconfigure, and test. Your data is not only in your data center but also likely in one or more cloud services. Therefore, you must make sure your DLP system extends to your cloud assets and is enforced. All cloud systems should be considered base components of your data repositories. Because most DLP systems are centrally managed and contain central policies, a routine process should include reviewing and testing each rule. DLP policies can reach down into the endpoint to ensure unauthorized USB thumb drives and other types of unapproved media are disallowed and alerted on. If your DLP solution has these or other capabilities, you should review each item and enforce where corporate policy allows.

TIP For more information, consult the NIST SP800-171 Rev. 2, Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations, at https://csrc.nist.gov/publications/detail/sp/800-171/rev-2/final.

Content Filter/URL Filter

Web content filtering is the practice of blocking access to web content that may be deemed offensive or inappropriate, or even contain dangerous items. Often after an investigation you will find that the attacker used a URL or keyword that was not properly blocked, sanitized, or removed from the email or web page; that may have led to the exposure or compromise. Numerous no-cost external third-party lists can help start your base list. Most modern content filter systems should have a subscribe to list that consistently updates the system automatically for questionable URLs and content. On top of that you should add/update the lists to ensure you are blocking sites involved in the compromise. Figure 29-4 shows how a content filter works.

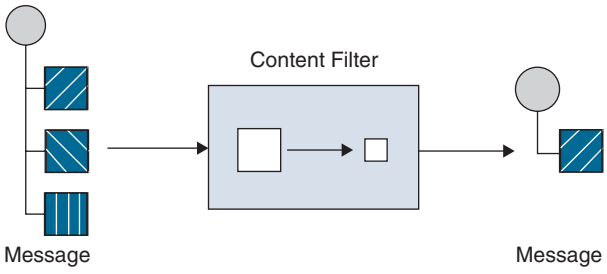


FIGURE 29-4 Content Filter Mechanics

TIP Although NIST doesn't have a specific SP dedicated to URL and content filtering, you can find references and best practices in various guides such as NIST SP 800-95 and NIST SP 800-53.

Key Topic

Update or Revoke Certificates

Certificates are among the most important methods for authorizing device access to your network. When one of these systems or hosts has been compromised, lost, stolen, or taken over, it is important to act quickly to remove or revoke the certificate authorizing access. Certificate revocation is the act of invalidating TLS/SSL certificate access before its scheduled expiration date. A certificate should be revoked immediately when its private key shows signs of being compromised. To revoke a certificate from a certificate authority (CA), each system has specific tasks/requirements; in Microsoft, the MMC snap-in allows you to right-click on it in the Issued list, go to All Tasks, and then choose Revoke Certificate. The interface then asks you for a reason code and a timestamp. You enter the data as required and then select OK.

Figure 29-5 shows how a client's web browser checks the status of the certificate by checking a certificate revocation list (CRL) of the CA.

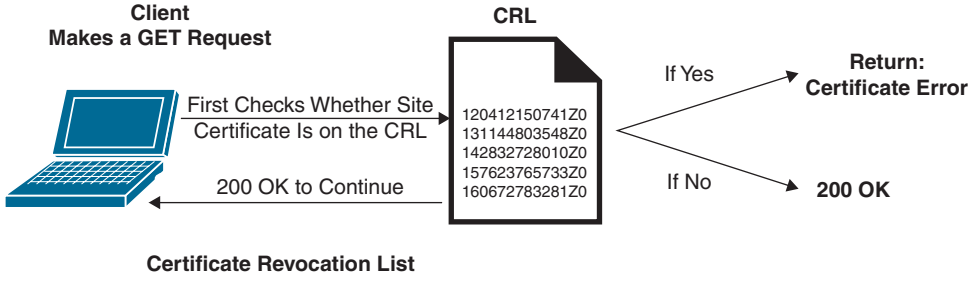


FIGURE 29-5 CRL Checking Process

TIP RFC5280 details x.509 public key infrastructure certificates and certificate revocation lists. You can find it at https://tools.ietf.org/html/rfc5280.

Isolation

Once a compromised endpoint has been isolated because of malware, a virus, or a Trojan, or it is part of a much wider attack, it is automatically quarantined and isolated or manually isolated. At this point, you should immediately consider three things:

1. Is there something on this host (private or personal data, or critical evidence) that cannot be obtained from some other source?
2. Can you safely reimage the device? In other words, can you trust underlying systems, BIOS/UEFI, drivers, firmware, or other operating systems after the reimage?
3. Can you wait until the investigation is completed before attempting any action?

These questions and corporate policy will help you determine the right course of action. Isolated hosts are good sources of evidence either on the disk/media or in memory, so investigators often want them intact. Investigators can build a list of what was done and how. That list then can be used to help protect these systems in the future. There are ways to isolate a host from the rest of the systems and network so that it can only operate in a sandbox to be observed by analysts and investigators during an incident; this approach allows for greater evidence collection.

TIP NIST SP 800-123 and SP 800-83 are general guides to server and host security best practices including compromised device ***isolation***. See https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-83r1.pdf.

Containment

When you know where the issue is and what systems have been affected, you should isolate or quarantine those systems so that threats cannot spread through the network. To do this, you need to disable network access for the affected computers and devices, or place them in a sandbox network. Then you need to reset passwords and

other credentials, like certificates, and block accounts that seem suspicious or that are not on the list of approved accounts with access. Make sure you don't delete information that may be useful in forensic investigations, which could help ensure this kind of event does not occur again. After the system has been released, you should run the appropriate measures to recover, remediate, or reimage the system. Ultimately, for a system to have integrity, it is highly recommended that system firmware and BIOS/UEFI be checked and verified and then the system be reimaged.

To perform ***containment*** on a system, you disconnect the network connection or disable the port on the switch. If you're using the system for forensic examination, you create a sandbox network that has faux services like DNS and other hosts to monitor the infected or suspect host.

Segmentation

Key Topic

Network ***segmentation*** is an architectural approach that divides a network into multiple networks, subnets, or segments. This segmentation allows network administrators to control the flow of traffic between these networks based on a granular policy. After an incident/investigation has been completed and things return to normal operations, network administrators and engineers should segment high-value networks.

Another type of isolation is more of a secure asset protection segmentation method. In this method, you can use private virtual local-area networks (VLANs) and access control lists (ACLs) to allow only certain hosts in certain networks to talk to other specific hosts, ensuring that device is protected from unwanted access. The inherent weakness of assumed trust is the key reason that many organizations have begun to adopt the zero trust strategy. Zero trust assumes nothing is trustworthy by default, even those already inside the network perimeter.

There are a few types of network segmentation:

- Physical segmentation, or separating networks via a firewall or other access controls
- Virtual segmentation, where VLANs divide network traffic but devices can still communicate with each other

Figure 29-6 shows network segmentation using firewalls to separate networks. In the middle is a bastion host that both sides can access, typically for sharing files and exchanging data.

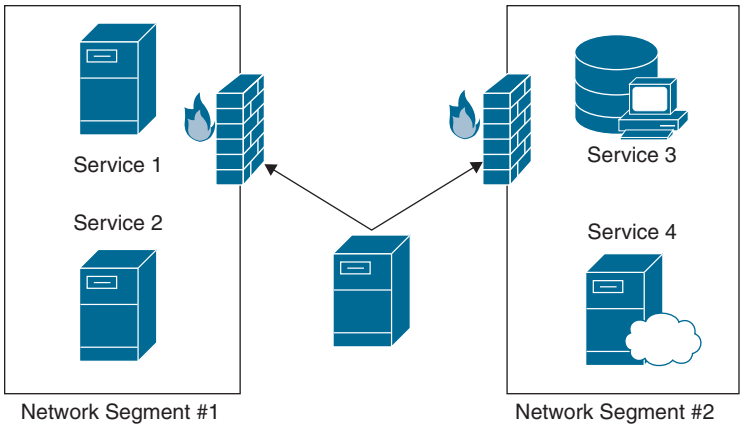


FIGURE 29-6 Network Segmentation

Segmenting a network leads to better performance due to less congestion with fewer devices. With fewer hosts per subnet, a segmented network minimizes local traffic and reduces the "noise" in broadcast traffic. Another benefit of a network segmentation policy is that it helps with identification and compliance requirements. One important item to remember after an attack is to split a network into zones that contain data with the same compliance or regulatory rules and restrict access to only authorized IP source addresses.

Key Topic

SOAR

Security Orchestration, Automation, and Response (SOAR) alerts for suspected phishing emails, endpoint attacks, failed user logins, malware, and other threat information that come from a variety of detection sources, such as SIEMs, systems, switches, and logging services. As the SOAR platform aggregates the suspected threats, it automatically triggers a process to inform affected users about the possible threats.

As part of the triage process, the SOAR platform is capable of extracting indicators of compromise (IoCs). The SOAR platform assigns an incident severity value and checks for cross-referencing the data against external threat intelligence databases. If any malicious indicators are found, affected users and administrators are informed with instructions on what to do. The SOAR's main purpose is to take routine actions and then automate them where a common issue exists, thereby reducing the involvement of an analyst. After an incident, thoroughly reviewing the threat and then building a playbook to respond to that specific item will reduce time to resolution and improve your security posture.

Key Topic

Runbooks

A ***runbook*** consists of a series of conditional steps to perform actions such as enriching data, containing threats, and sending notifications automatically as part of the incident response or security operations process. This automation helps accelerate the assessment, investigation, and containment of threats to speed up the overall incident response process. Runbooks can also include human decision-making elements as required, depending on the particular steps needed within the process and the amount of automation the organization is comfortable using. Like playbooks (described in the next section), runbooks can also be used to automatically assign tasks that will be carried out by a human analyst; however, most runbooks are primarily action-based. Figure 29-7 shows a typical automated runbook process flow. When building runbooks, you select components and tie them together to create the desired process flow. Here you can assume a file will be checked as a virus, so you start with the "Start" process. Next, the DNS service does a reverse lookup, potentially obtaining the validity of the file. After that, two processes run: a hash is taken and compared, and the file is submitted to Virus Total to see whether it has previously been reported as a virus. Runbooks can also provide step-by-step processes that a person should follow.

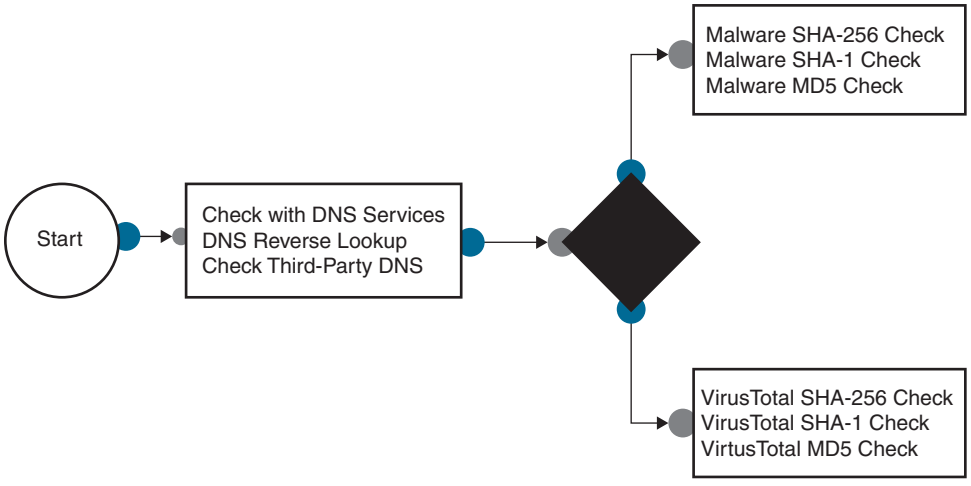


FIGURE 29-7 Automated Runbook Flow

NOTE IT operations use a runbook for reference for routine procedures that administrators perform. A playbook normally provides manual orchestration of incident response; newer systems enable you to build and automate runbooks.

Key Topic

Playbooks

A SOAR ***playbook*** is a linear-style electronic checklist of required steps and actions needed to successfully respond to specific incident types or threats. Incident response playbooks provide a simple top-down step-by-step approach to orchestration. They help establish formalized incident response processes and procedures within investigations by identifying required steps systematically followed, which can help meet and comply with regulatory requirements. Playbooks support both human and automated actions. An effective use of SOAR playbooks is to document processes and procedures that rely heavily on tasks that humans carry out manually, such as breach notification or highly technical processes such as adding IP addresses to firewalls.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 29-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 29-2 Key Topics for Chapter 29

Key Topic Element	Description	Page Number
Section	Application Approved Lists	822
Section	Application Blocklist/Deny List	822
Section	Quarantine	823
Section	Configuration Changes	824
Section	Firewall Rules	825
Section	MDM	825
Section	Data Loss Prevention	828
Section	Content Filter/URL Filter	828
Section	Update or Revoke Certificates	829
Section	Isolation	830
Section	Containment	830

Key Topic Element	Description	Page Number
Section	Segmentation	831
Section	SOAR	832
Section	Runbooks	833
Section	Playbooks	834

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

application approved lists, application block list, deny list, quarantine, mobile device management (MDM), data loss prevention (DLP), isolation, containment, segmentation, Security Orchestration, Automation, and Response (SOAR), runbook, playbook

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What is the purpose of an application approved list?
2. When a file or application is quarantined, what happens to it?
3. In data loss prevention (DLP) systems, how is exfiltration of sensitive data protected?
4. What is the purpose of revoking a certificate?
5. SOAR requires runbooks and playbooks. Whereas playbooks consist of a number of plays, runbooks are a series of what types of steps?



This chapter covers the following topics related to Objective 4.5 (Explain the key aspects of digital forensics) of the CompTIA Security+ SY0-601 certification exam:

- Documentation/evidence
  - Legal hold
  - Video
  - Admissibility
  - Chain of custody
- Timelines of sequence of events
  - Time stamps
  - Time offset
  - Tags
  - Reports
  - Event logs
  - Interviews
- Acquisition
  - Order of volatility
  - Disk
  - Random-access memory (RAM)
  - Swap/pagefile
  - OS
  - Device
  - Firmware
  - Snapshot
  - Cache
  - Network
  - Artifacts
- On-premises vs. cloud
  - Right-to-audit clauses
  - Regulatory/jurisdiction
  - Data breach notification laws
- Integrity
  - Hashing
  - Checksum
  - Provenance
- Preservation
- E-discovery
- Data recovery
- Non-repudiation
- Strategic intelligence/ counterintelligence

Understanding the Key Aspects of Digital Forensics

Digital forensic science is a branch of forensic science that is focused on the recovery and investigative materials found in digital devices related to cybercrime. Digital forensics is the process of identifying, preserving, analyzing, and documenting digital evidence. This is done in order to solve a crime and present evidence in a court of law. There are five steps to digital forensics:

1. Identification
2. Preservation
3. Analysis
4. Documentation
5. Presentation

Digital forensics also has nine phases, which are covered in this chapter.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 30-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 30-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Documentation/evidence	1–4
Acquisition	5–7
Integrity	8
On-premises vs. Cloud	9

Foundation Topics Section	Questions
Preservation	10
E-discovery	11
Data Recovery	12
Nonrepudiation	13
Strategic Intelligence/Counterintelligence	14

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is the first step when litigation has been filed or will soon be filed, with a focus on saving data in its current state?
   a. Data recovery
   b. Data collection
   c. Data destruction
   d. Data preservation

2. Which of the following is true regarding the correct format for forensic video evidence?
   a. Video settings should always match the video recorder; otherwise, the playback will be corrupted.
   b. Video formats and settings can vary.
   c. Videos are typically recorded at the same FPS and resolution.
   d. The frames per second (FPS) and resolution cannot be changed.

3. Chain of custody in forensic investigations requires investigators to perform which of the following?
   a. Tag some evidence so that it can be chronologically stored in a locked storage facility.
   b. Clearly and accurately document the collection process, methods, and path of evidence from collection to presentation.

c. Store evidence in humidity-free lockers to ensure integrity and chain of custody.

d. Take responsibility for all evidence collection after the site has been turned over from the analyst.

4. Which of following describes the use of a tag in forensic investigations?

a. Tags provide a history identifying when items sold.

b. Tags can contain date/time and investigator initials to help determine who collected items and when.

c. Tags allow forensic investigators to sell the items later after the investigation has been completed.

d. Tags help investigators determine whether the evidence was collected onsite or during the reverse-engineering process.

5. When you are collecting evidence, the order of volatility requires you to do what?

a. Start evidence collection with the most volatile data/information.

b. Start evidence collection after the volatile data has been removed from the site.

c. Start with the least volatile evidence to make sure your evidence is intact when received at the office.

d. Ensure that the most volatile data/information that is collected is handled by the police only.

6. Which artifacts can be collected during a forensic investigation of a suspected workstation that will help or contribute to the validity of the investigation?

a. Data and evidence stored in appropriate containers

b. Art, photos, and images that contain facts about the attack and that can be used to piece together a picture of what took place

c. Log files collected from previous years, log file configuration settings for global policies, and SCCM data

d. All logs, the registry, the RDP cache, and Windows Error Reporting (WER)

7. What volatile memory item contains more trusted information than disk, magnetic, or optical storage?
   a. ROM
   b. RAM
   c. RTU
   d. CPU

8. When dealing with a forensic investigation, you must take images of suspect hard drives. Which of the following best describes the process to ensure a hard drive hasn't been tampered with after the collection/image has been completed?
   a. Checksums provide a thin piece of paper that acts as a seal over the evidence to ensure no one opens that checked evidence.
   b. Checksums provide checks and balances during an investigation to ensure investigators are working from the most important data to the least important.
   c. Checksums provide a check and sum for each bit of data collected, and when they are combined with a hash, you can track data.
   d. Checksums of a newly imaged drive/data can be created, so if the device/data is modified, the checksum will not match, letting you know it was modified.

9. Which statement regarding on-premises and cloud-based systems is true?
   a. Maintaining a chain of custody isn't very challenging in a cloud environment versus a traditional forensics environment.
   b. *On-premises* and *cloud-based* are simply terms describing where data is stored.
   c. When a cyber incident happens, legal jurisdiction and the laws that govern the region do not normally present unique challenges.
   d. In traditional forensics environments, the external security team has control over who is conducting forensics operations on a machine.

10. What happens when digital evidence is not properly preserved and collected in accordance with forensic best practices?
    a. It becomes volatile.
    b. It is known as nonprovenant.
    c. It becomes an artifact.
    d. It is not admissible in court.

11. What should organizations have to perform to preserve information during e-discovery?
    a. Legal hold process
    b. Nonrepudiation
    c. Order of nonvolatility
    d. Counterintelligence

12. Which of the following is the extraction of data from damaged, deleted, or purposely destroyed evidence sources in a forensically sound manner?
    a. E-discovery
    b. Admissibility
    c. Data recovery
    d. Strategic intelligence

13. Which term best defines the assurance that someone cannot deny the validity of something—where a statement's author cannot dispute its authorship?
    a. Snapshots
    b. Timestamps
    c. Checksums
    d. Nonrepudiation

14. Which of the following is information gathered and activities conducted to protect against espionage, other intelligence activities, or sabotage conducted by or on behalf of other elements?
    a. Counterintelligence
    b. Strategic intelligence
    c. Artificial intelligence
    d. Intellectual property (IP) intelligence

Foundation Topics

Documentation/Evidence

Digital forensics processes include documentation of evidence and reporting. In digital forensics, the goal is uncovering and interpreting electronic data. The process is to preserve any evidence in its original form while performing a structured investigation by collecting, identifying, validating, and clearly documenting the digital information to reconstruct past events. Careful and meticulous records of all the data to re-create the crime scene must be taken in order for a forensic case to be successfully presented and defended in court.

Legal Hold

Key Topic

A ***legal hold*** plan is essential guidance for information governance or an e-discovery strategy, recognizing when the duty to preserve begins, what it entails, and how to implement it and when it ends. A legal hold is a legal document used to communicate to a party or custodian that litigation is imminent or reasonably anticipated, potentially relevant information must be preserved, and electronically stored information may be relevant to pending litigation. Data preservation is the first step when litigation has been filed or will be soon filed, with a focus on the preservation of data in its current state, such as emails, SMS, MMS, and deleted messages (still on disk and not destroyed) on all devices including cell phones, PCs, and mobile devices. All organizations should have a corporate policy in place when dealing with third parties, which gives them a legal hold over data.

Video

Key Topic

Forensic *video* analysis is the scientific examination, comparison, and/or evaluation of video in legal matters. Forensic video expertise is essential with the proliferation of smartphones and security cameras. Today more incidents are being captured on video than ever before. The increase in the amount of video evidence cuts both ways: a major incident can overwhelm law enforcement agencies with the sheer amount of video that needs to be stored, processed, and reviewed.

Beyond the vast amount of video now available, digital video is much more complex in not only the images but also the data it contains. In most cases if video does not have proper forensic analysis, the courts will not accept it at its face value. There are a large number of factors to consider, from the method video is recorded to the color calibration of the monitor. All of these factors can have a significant impact on what people see and how they interpret the video. These factors can easily lead to incorrect conclusions about what really happened. The way video and images are

captured can have a dramatic impact on how people interpret what they see, which can also lead to incorrect conclusions.

Many different camera makes and models are on the market, including dedicated digital photography cameras, action cameras, smartphones, and security system cameras. File formats can vary, as well as settings for recording video, such as frames per second (FPS) and video resolution. These features can all factor into how and what video information is stored. A forensic video analyst must understand all of these issues to prevent a misinterpretation of what a video appears to show, such as colors or perhaps the speed of events unfolding in the video. In a criminal case, where a video could be very important to proving someone's innocence or guilt, settings and other factors can be very important. Even now with deep fakes and artificial intelligence (AI), it's more important than ever to ensure the proper investigative processes, procedures, and handling of video evidence.

Key Topic

Admissibility

Digital evidence is *admissible* if it establishes a fact-of-matter assertion in the case, and it must remain unaltered during the digital forensics process. The results of the examination must be valid, reliable, and peer reviewed. Forensically sound methods and tools must be used to obtain the digital evidence, and the authenticity, integrity, and reliability of this evidence must be supported by the expert witnesses and digital forensic analysts who processed it.

Three standards of evidence must always be considered:

- Is it sufficient (which is to say, convincing without question)?
- Is it competent (which means it is legally qualified)?
- Is it relevant (which means it must matter to the case at hand)?

To be admissible, the findings should be interpreted in an unbiased manner, and errors and uncertainties in the findings, as well as limitations in the interpretations of results, should be disclosed.

Whether evidence is admissible is determined by following three rules:

- Best evidence means that courts prefer original evidence to copies to avoid alteration of evidence.
- The exclusionary rule means that data collected in violation of the Fourth Amendment (no unreasonable searches or seizures) is not admissible.
- Hearsay is second-hand evidence and is often not admissible, although some exceptions apply.

Chain of Custody

Key Topic

For evidence to be credible in court, you must follow strict rules. Following the proper ***chain of custody*** consistently and methodically has been challenging due to the dynamic nature of digital evidence. Cyber attacks and cybercrimes are evolving, and the invincible nature of these attacks makes it difficult to gather evidence. Investigators are required to adapt to changing crime scenes and the digital media that has been used to commit these crimes. It has become imperative that standard procedures are coherent and ensure harmony between the legal parties involved.

NOTE The chain of custody provides a clear record of the path that evidence takes from acquisition to disposal. It is often required in court proceedings to prove evidence isn't tampered with.

Timelines of Sequence of Events

Key Topic

The ultimate goal of timelines in digital forensics is to reconstruct what happened, presenting a narrative about an event that is coherent, believable, and supported by sufficient evidence. Timelines have become a mainstay of digital forensic analysis in both public and private sectors. They help explain what was happening on a given device or set of devices during a cybersecurity incident, crime, or other event. The timeline helps frame the situation to explain it to attorneys, juries, and other stakeholders.

Timestamps

Timestamps play a very important role in many digital forensic examinations, so it's very important for any forensic examiner or analyst to clearly understand how they work. To prove that a certain document was created before or at a certain time, you can timestamp it.

An examiner must know the ins and outs of each operating system; for example, to know that NTFS has eight timestamp values called Modified, Accessed, Changed, and Birth (MACB) times, of which $STANDAR_INFO and $FILE_NAME both contain a set of four each. Certain timestamps can be modified by users, so knowing which files can be changed and which reflect the right date/time is important to an investigation.

To view them the investigator runs **MFTRCRD.exe c:\file\file.a -d indxdump=off 1024 -s**.

Time Offset

In a forensic examination, establishing the time zone from the suspect system is one of the first tasks for a forensic examiner. If the offset information is not established at an early stage and taken into account, the validity of date/time evidence may be brought into question. This is true for the examination of browser history and related artifacts, as well as the examination of file system metadata. Some date/time values stored in binary files are affected by the time zone setting of the original computer, and many digital forensic applications can alter the representation of these dates by the time zone setting of the forensic workstation. The situation becomes particularly complicated when the suspect computer was set to an incorrect time zone and the computer clock was set to correspond to the local time zone. Many date/timestamps store data as Coordinated Universal Time (UTC) values. If this is the case, the operating system and some applications have to convert the value from local time to UTC. On Windows systems, you can check the system registry at HKEY_LOCAL_MACHINE\System\Select, where four keys detail the Current, Default, Failed, and LastKnownGood control sets. UTC = LOCAL TIME + BIAS, where active time bias is the current time difference from UTC in minutes.

NOTE Time offset is recorded against a verified time standard. Recording the time offset is critical to an accurate examination when dates and times are at issue. It should be recorded at the beginning of any examination.

Key Topic

Tags

The actual collection and processing of evidence are critical steps in the investigative process. Each piece of evidence collected must be handled in a way that preserves its integrity and any trace evidence. The collection process must provide a detailed record of an item's whereabouts from the time of collection to the time it arrives in a courtroom. Failure to pay proper attention to any one of these areas can easily result in one or more pieces of evidence having no value in court or in administrative proceedings. After an object is identified as evidence, it must be tagged. Evidence tagging helps identify the collected item. The ***tag*** can consist of something as little as a sticker with the date, time, control number, and name or initials of the investigator. Using a control number is an easy way to identify a piece of evidence in documentation such as a chain of custody. A tag can also be an actual document that contains general information about the item and the incident under investigation.

After the evidence is tagged, the investigator should photograph it in a way that also displays the tag information. Photographing evidence becomes another way to document what was collected and how it was processed. When taking pictures of computing devices, you should include all interfaces. If a cable is attached to an

interface, it should remain connected during the picture-taking process. It's a good practice to clearly label each attached cable with the associated peripheral device before taking interface photos. You should place all items in protective bagging. For electronic devices, you should use a Faraday/antistatic bag to ensure no outside radio frequency (RF) can influence the device.

Reports

Every step in the investigation needs to be reported on. Many reports can make up the final report provided as a summation of the entire case. Each report offers its own insights into intent and activity, and the tools deployed are always directly tied to the case at hand and specific hurdles that are faced. As with all evidence, well-documented reports ensure that no corners are being cut and each aspect of the process is forensically defensible. Investigators can use existing templates to guide them along the path of process, procedures, controls, and requirements.

Event Logs

On Windows systems, ***event logs*** contain a lot of useful information about the system and its users. Many analysts rely on Windows event logs to help gain the context of attacker activity on a system. Unfortunately, some things cannot be determined using only the event log. Collecting Windows file activity is a massive undertaking. Locating and storing logs from the system should be your first step in evidence collection. For example, Windows Event Viewer logs are stored in %SystemRoot%\System32\Winevt\Logs. Each system contains these critical and useful files in different places depending on the operating system. Based on the system you are investigating, you need to obtain the appropriate location for these files (prior to touching the device).

TIP You should know the various Windows Event Viewer logs and what they do. The Windows Event Viewer security log provides information related to the success and failure of login attempts as well information related to other audited events. The application log provides information related to applications run on the local system. The system log records error messages and other information generated from the Windows operating system itself.

Interviews

Different types of crimes such as computer crime, fraud, and hacking require different interview methods. These types of crimes can be performed using

computers or mobile devices, including smartphones, tablets, and digital cameras. To have a successful interview, the interviewer should collect the following information:

- Background details regarding the victim or device
- The focus of the investigation
- The possible perpetrators and their background information

During interviews, interviewers should try to recognize new details and other witnesses who could help in resolving the case. Different interview techniques exist, but interviews usually should try to answer simple questions such as who, when, where, what, how, and why. The initial interview is typically the best chance to collect basic evidence from perpetrators who used computers and mobile devices as well as information from victims. After information necessary for the interview is gathered, the interview process can start. At this point, the following must be ready:

- Privacy Act statement
- List of official papers from the interviewee
- Checklist with information gathered prior to the interview
- List of questions
- Copies of all official papers planned to show to the perpetrator or victim
- The method of recording the interview

The interview should be conducted in a peaceful and comfortable setting. Considering the interviewees as individuals and cooperating with them in a serious and cautious manner allows the interviewer to develop a rapport. A good relationship during the interview between the investigator and the victims or perpetrators is important to achieving good results. The interviewer should always try to create common ground to make interviewees relaxed and comfortable. At the same time, the interviewer should avoid a heavy-handed approach that enforces authority. An authoritative approach yields less successful results because it makes victims or perpetrators less comfortable and on the defensive.

Acquisition

Forensic ***acquisition*** is the process of collecting specific data related to an attack, intrusion, or investigation. Items can include computer media such as hard drives, thumb drives, servers, phones, tablets, and other devices that store electronic data. Investigators need to follow clear and precise steps when acquiring data from a crime/investigative scene so that all collected data retains its integrity.

Order of Volatility

The collection of evidence should start with the most volatile item and end with the least volatile. The ***order of volatility*** is the order in which the digital evidence is collected. Highly volatile data resides mostly in the memory, cache, or CPU registers, and it will be lost as soon as the power to the computer is turned off. Less volatile data is a not lost so easily and is relatively stable because it may be stored on disk drives or other permanent storage media, such as CD-ROM discs. Crime scene technicians should collect evidence beginning with the most volatile and then move toward the least volatile. The following order of volatility was taken from RFC 3227, Guidelines for Evidence Collection and Archiving:

1. Cache, registers
2. ARP cache, routing table, memory, kernel statistics, process table
3. Temporary files
4. Disk
5. Monitoring data and remote logging pertaining to the computer in question
6. Physical configurations, network topology
7. Archival media

Disk

The hard drive or solid-state drive (SSD) is usually the center of most cyber forensic investigations. A forensic image is an electronic copy of a drive. It's a bit-by-bit or bitstream file that's an exact, unaltered copy of the media being duplicated. Even though you might think that the data you place on a disk will be around forever, that is not always the case. The likelihood that data on a disk cannot be extracted is very low. The first step in hard drive forensics is the identification of storage devices at the scene.

Forensic disk controllers are most commonly associated with the process of creating a disk image, or acquisition, during forensic analysis. Their use is to prevent inadvertent modification of evidence. Disk drives are considered a nonvolatile form of data storage.

Random-Access Memory

There are several reasons that a complete random-access memory (RAM) capture might prove useful; most often they revolve around the key differences between data stored in RAM and data stored on a hard disk drive. RAM is considered volatile

memory and is perceived to be more trusted than nonvolatile memory, such as ROM, disk, magnetic, or optical storage. Investigations using live forensic techniques require special handling because the volatile data in RAM can be lost if the system is turned off.

If we consider data that is either not stored or somehow protected on a hard disk, yet stored in plaintext when stored in RAM, many data types immediately come to mind: passwords, financial transaction information, encryption keys, and similar types of high-value data. The type of information in RAM might not even be intentional (poorly written apps can leave residual data), or it could be a byproduct of malware. Certain malware can be completely memory resident and leave critical information available that may lead to the author. RAM contains

- Unsaved documents
- Passwords
- Credentials
- Who is logged in to the system

Code from programs that are not necessarily written to the computer and/or saved, but resident, can also include printed pictures, emails, chat messages, malware, running processes, and more. During every second of the computer's use, RAM is changing, writing, or reading something that the computer is actively working on. It may now become apparent why collecting RAM is the priority in a live forensic triage or on-scene digital investigation. To dump the volatile memory, options are available for all platforms. On Windows, you can perform a kernel or a full dump directly or use a third-party tool specific to your platform. Due to operating system security restrictions, forensic programs occasionally are not entirely executed in the user space; therefore, you must instead inject a special kernel driver or module into the core of the operating system to obtain dumps of memory.

Key Topic

Swap/Pagefile

A ***swap file*** or ***pagefile*** is an area of the hard drive that the operating system can use to store data from RAM that has not been used in recent activity, typically moving data/apps required from disk to RAM and from RAM back to disk, hence the name "swapping."

Pagefile.sys is used within Windows operating systems to store data from RAM when it becomes full. Pagefile.sys is a contiguous file, so it can be read more quickly; it is located on the root of the hard drive; and normally, the more infrequently used files are stored to it. It can also be used as a data backup in the event of a system crash. By default, the Windows operating system configures the size of Pagefile.sys;

however, it can also be altered by the user. Normally, Pagefile.sys can be a significant proportion of data present on the hard drive, but removing it can greatly reduce the operating speed of the computer.

The swap file or pagefile is also called *virtual memory*. During a forensics investigation, a pagefile is very important and should be treated as such. Although not as volatile as RAM, it should always be inspected using forensic tools because it might reveal critical details such as passwords and other important forensic artifacts.



Operating System

Operating system forensics is the process of retrieving useful information from the OS of the computer, PC, or mobile device in question. OS forensics involves forensic examination of the PC's operating system. The foremost commonly used operating systems are Windows, macOS, and Linux. The file system also identifies how disk drives store data. There are many file systems introduced for various operating systems, like FAT, exFAT, and NTFS for Windows operating systems, and Ext2fs or Ext3fs for Linux operating systems. Data and file recovery techniques for these file systems include data carving, slack space, and data hiding.

Another important aspect of OS forensics is memory forensics, which includes virtual storage, Windows memory, Linux memory, macOS memory, memory extraction, and swap spaces. OS forensics also involves web browsing artifacts, like messaging and email artifacts. Some indispensable aspects of OS forensics are discussed in subsequent sections.

Another option for forensic analysts in identifying, collecting, and preserving evidence from computers is to use a bootable forensic toolset/disc, one that has a special operating system and tools for ensuring that the original system is not modified. The SIFT workstation, Sleuth Kit, Autopsy, FTK (Forensic Toolkit), Kali forensics, and many others can help with your investigation.



Device

Depending on the forensic investigation and the device(s) involved, a multitude of devices, software, scripts, and systems can be used. From cell phones, tablets, PCs, networking devices, printers, and many other computer-based devices (like desk phones and voice assistant devices), forensic examination can be performed on these devices and many others. You should choose the appropriate tool and methods for the device and expected results. For example, when you're performing forensics on a mobile platform device like a cell phone, there are several methods and many tools. They range from automated, physical, and logical to brute-force acquisition tools like svStrike, IP-Box, FTK Imager, Zune, and Xways. Two of the most common tools for cell phone forensics are Cellebrite and EnCase. As previously mentioned,

you need to use the right platform, the right software, the right process, and the right procedures for the specific device you are investigating.

Key Topic

Firmware

System firmware like the basic input/output system (BIOS) or Unified Extensible Firmware Interface (UEFI) on more modern systems is the first program that runs on the CPU when a computer is turned on. System firmware has a much more privileged software layer in computer systems and more recently has become the target in sophisticated computer attacks. Malware used by high-profile rootkits is almost completely invisible to standard forensic procedures and can be detected only with special mechanisms. Researchers have provided insights into well-known open-source memory forensic tools and have evaluated the approach within both physical and virtual environments. The firmware code initially runs on the ROM chip. Then it moves from ROM into RAM by manipulating special registers, called Programmable Attribute Maps (PAMs), to shadow the ROM-mapped regions with RAM and uncompressing its code image into memory. Firmware then starts initializing devices on the system bus and maps their registers and memory into the physical address space as required.

The pre-boot phase, which is the actual power-on self-test (POST), loads firmware settings, checks for a valid disk system, and if the system is good, moves to the next phase, verifying that the computer has a valid master boot record (MBR) before loading the Windows boot manager. As you can see, firmware all along this path is referenced and loaded multiple times, giving rise to the potential for an attacker's interaction with this load-in process. A report from *Wired* magazine shows that many BIOSs share portions of the same code, which leads to incursion vulnerability levels in the 80 percent range. This code reuse is even present in big-brand PC makers. Attackers are looking to modify and sabotage firmware to target specific sections of the operating system, and infiltrate other running software. Securing firmware is largely in the hands of the product designers' hands. Intel's Hardware Shield, Microsoft's updated OS protection, and Dell's enhanced BIOS verifications are all evidence of vendors taking responsibility. You too can help by updating your firmware when an update is released. If a USB drive is required, make sure it is one you can trust. If you're purchasing new hardware, ensure it has advanced firmware security.

Key Topic

Snapshot

A snapshot preserves the entire state and data of the virtual machine at the time it is taken. A snapshot includes the virtual machine settings and the state of the machine's virtual drives. It can also be used as a restore point when testing software or configuration changes fail.

Because computers are vulnerable to cyber attacks, cybercriminals can counterfeit and fabricate evidence stored on the computer in question. Forensic analysts therefore must protect the evidence from loss. Before an investigator starts examining the computer for digital evidence, the entire drive of the computer should be imaged (copied) to preserve data and verify its integrity. In various circumstances, crime scene investigators create a bitstream copy of a storage device with the help of a forensic imaging tool. Then they store that original storage media on a forensically clean storage device and subsequently work on the image obtained.

Proper checks are necessary for successful image duplication, requiring the image makers to perform a hash calculation before and after the creation of a forensic image. A hash calculation verifies that the image wasn't altered or damaged during an imaging process. If the duplication is successful, the hash of both the original copy and imaged copy should be the same.

Cache

Key Topic

The contents of the CPU cache and registers are extremely volatile because they are changing all the time. Nanoseconds literally make the difference here. An examiner needs to get to the cache and register immediately and extract that evidence before it is lost or corrupt. Users utilize web browsers for many functions, such as searching for information, accessing email, engaging in e-commerce or banking, instant messaging, visiting online blogs, and accessing social networks. Web browsers record many data-related user activities; information such as URLs visited by users, search terms, cookies, cache files and images, access time, and use time is held in memory on the system. Examining the evidence that is the subject of criminal records is an important step in the examination of your browser. Being able to uncover an offender's profile and connections depends on the web registry. The evidence obtained from the use of the web browser is a key component of a forensic expert's case. It is possible to analyze evidence such as history, cache, cookies, downloads, URL addresses, access times, and frequency of visits from a suspect's computer. The analysis process should first identify the web browsers.

Network

Key Topic

Network traffic and logs can provide empirical evidence if forensic analysts properly collect and preserve them. Various network devices and tools, including switches, routers, VPN appliances such as SIEMs, proxies, and firewalls, can be configured to record logs of the activities and events that occur on them and can provide evidence for court proceedings. Network forensic analysis tools (NFATs) typically provide the same functionality as packet sniffers, protocol analyzers, and SIEM software, sometimes in a single product. Whereas SIEM software concentrates on correlating events among existing data sources that typically include multiple network traffic

and related sources, NFAT focuses primarily on collecting, examining, and analyzing network traffic. NFAT software also offers additional features that further facilitate network forensics, such as reconstructing events by replaying network traffic within the tool, ranging from an individual session like instant messaging (IM) between two users to all sessions during a particular time period. The speed of the replay can typically be adjusted as needed.

Key Topic

Artifacts

Forensic ***artifacts*** are remnants of an intrusion that can be identified on a host or network. Artifacts can be extracted from hosts such as logs, the registry, browsers, the RDP cache, and Windows Error Reporting (WER). Artifacts that can be extracted from hosts within the network or the compromised network itself lead to the entry point of a threat actor and the exfiltration point(s).

Common examples of indicators of compromise (IOCs) include

- IP addresses
- IPv4 and IPv6
- URL and FQDN+Path
- MD5 hash and SHA-1
- Hash
- Filename and type
- Windows registry key
- Windows driver

On-premises vs. Cloud

Cloud computing has transformed the IT industry; services can now be deployed in a fraction of the time that they used to take. Scalable computing solutions have spawned large cloud computing companies such as Amazon Web Services (AWS), Google Cloud, and Microsoft Azure. With a click of a button, personnel can create or reset the entire infrastructure of a computing resource in three main types of cloud computer service models: software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS). These models present three unique challenges to conducting cloud forensic investigations.

Forensic issues that are unique to cloud computing are jurisdiction, multitenancy, and dependency on cloud service providers (CSPs). Cloud forensics is a subset of digital forensics based on the unique approach to investigating cloud environments.

CSPs have servers around the world to host customer data. When a cyber incident happens, legal jurisdiction and the laws that govern the region present unique challenges. A court order issued in a jurisdiction where a data center resides is likely not applicable to the jurisdiction for a different host in another country. In modern CSP environments, the customer can choose the region in which the data will reside, and this decision should be made carefully.

A main concern for an investigator is to ensure that the digital evidence has not been tampered with by third parties so that it can be admissible in a court of law. In PaaS and SaaS service models, customers must depend on the cloud service providers for access to the logs because they do not have control over the hardware. In some cases, CSPs intentionally hide the details of logs from customers. In other cases, CSPs have policies that they do not offer services to collect logs.

Maintaining a chain of custody is very challenging in a cloud environment versus a traditional forensics environment. In traditional forensics environments, the internal security team has control over who conducts forensics operations on a machine, whereas in cloud forensics, the security team has no control over who the CSP chooses to gather evidence. If the investigator is not trained according to a forensic standard, the chain of custody may not hold up in a court of law.

TIP *Cloud-based* and *on-premises* are simply terms that describe where systems store data. Many of the same vulnerabilities that affect on-premises systems also affect cloud-based systems.

Key Topic

Right-to-Audit Clauses

In cloud-based systems, the ***right-to-audit*** provision of hosting, services, and data storage contracts is important, and it should be clearly defined that the buyer has the right to examine the books and records of the seller to assure that the seller has complied with the contract and that no employee of the buyer received any funds either directly or indirectly. Vendor audit clauses can help control fraud and abuse by affording a discovery device in a fraud examination. When the right to audit is exercised, fraud examiners or auditors are generally looking for fraud by vendors and violations of company ethics policies. For example, they might be looking for faulty or inferior quality of goods, short shipments, or high prices when the goods can be bought directly and/or cheaper from the same or another vendor, goods not delivered, kickbacks, gifts and gratuities to company employees, or conflicts of interest.

Regulatory/Jurisdiction

Key Topic

When data is stored, managed, and accessed in the cloud and users are in any number of countries, regulatory and jurisdictional authority might not be clearly delineated. Forensic regulatory and jurisdictional intervention involves data being in private data centers, offsite data centers, and cloud virtual environments where users have on-demand access to processing resources, data, infrastructure, and applications through the private connection or via the Internet. The liability for compliance shifts from solely your responsibility to a joint responsibility between you and the provider. When it comes to multi-jurisdiction engagement, whether for investigation or e-discovery, there will always be issues with data protection laws in various locations. Different countries have different laws when it comes to data protection or privacy; for example, the EU is highly regulated by the General Data Protection Regulation (GDPR).

GDPR is obviously relevant to any organizations that have offices in Europe, and it goes well beyond that. For example, if your clients are storing any data for any EU residents, that organization would be subject to GDPR as well. When you're engaging in any data collection, it is important to discern whether your organization needs to comply with GDPR or other country-specific laws and regulations, and if so, what steps need to be taken to ensure that the only data collected is relevant to the matter as opposed to everything being collected. Generally, forensic teams overcollect to some extent to ensure that important documents are not missed, but with privacy changes, it is important to understand the laws and collect only what is relevant. Certain country jurisdictions may not allow you to collect data at all based on where you are located. Therefore, you must check with all local laws, regulations, restrictions, and privacy rights before collecting data. Otherwise, fines could reach millions of dollars. In contrast, when data is stored locally or in your on-premises owned-and-managed data center, you are responsible for maintaining compliance. Keeping clear records of your team's part in the process is essential in proving compliance during an audit.

Data Breach Notification Laws

Key Topic

On-premises breaches are handled differently from cloud-based breaches. In some cases, you might not be made aware of a cloud breach in the regulatory stated time for notification and thereby be responsible for fines. This is an important factor in compliance requirements when working with third-party cloud hosting. Another important factor is that you are responsible for initiating the forensic investigation for onsite breaches. Depending on the cloud vendor and the agreement, you might have to wait until they perform the entire investigation before you get to the actual issue. In some cases, you can request a joint cloud investigation where you can send a team to the providers' data center to perform the investigation on behalf of both

of you. Just about every U.S. state has some sort of data breach laws. California SB 1386, breach notification legislation, has radically changed the nature and tenor of responding to data breaches by introducing mandatory notification requirements and opening the door to significant regulatory fines and civil damages. These laws have affected the way evidence in data breaches must be collected and treated.

The gathering of evidence has always been a significant step in analyzing the cause and extent of data breaches. However, the use of forensic evidence and methodologies, such as preserving data so that findings can be verified and authenticated in litigation, has become more complicated and important with the advent of cloud hosting over the last few years. Not all network breaches lead to unauthorized access to personally identifiable information (PII) or protected health information (PHI), so forensic analysis can remove the need for expensive notification requirements that can be detrimental to the company.

Another growing trend is for regulators to question the procedures used by organizations in determining the scope of a breach and the numbers of persons to be notified. There has been a surge in civil claims following data breaches, with plaintiffs' attorneys being ready to argue that inadvertent loss of data during the initial breach response spoiled relevant evidence, leading to negative presumptions against the breached organization.

Integrity

Forensic investigators play crucial roles within the legal system and are constantly under various pressures when performing analytical work, generating reports based on their analyses, or testifying to the content of these reports. Maintaining the scientific integrity of these actions is paramount to supporting a functional legal system and the practice of a good part of the science. There must be transparency and professionalism in the forensic field to ensure integrity and to reduce conflicts of interest. Besides the investigators' and analysts' integrity, the evidence must maintain its integrity by having a closely documented and controlled chain of custody. There must be a systematic, objective, scientific analysis of the incident, not only because this is the most reliable approach, but because a true scientific understanding of how an incident or failure happened is the first step to empowering others to prevent such incidents from happening in the future.



Hashing

A hash calculation is performed before and after the creation of a forensic image. It ensures that the image wasn't altered during the duplication process. In addition, it is essential to verify periodically that the hash of the image copy being used for forensic examination has not changed or been altered. Therefore, a hash calculation

confirms that the results drawn from the image copy would legally apply to the original source.

Hash functions have four defining properties that make them useful. Hash functions are

- **Collision resistant:** A collision occurs where multiple inputs are found to produce a common output or common hash value. Because potential inputs are infinite and the output is a fixed length, collisions are bound to occur.
- **Deterministic:** For any given input, a hash function must return the same value each and every time that input is processed.
- **Computationally efficient:** You can expect that a hash function will be computationally efficient or, in other words, speedy.
- **Pre-image resistant:** All hash functions must be "pre-image resistant." The hash function should not provide any clue about the size or content of the input.

Key Topic

Checksums

To produce a checksum, you run a program that puts that file through an algorithm. Typical algorithms used for this process include MD5, SHA-1, SHA-256, and SHA-512. The algorithm uses a cryptographic hash function that takes an input and produces a string, a sequence of numbers and letters of a fixed length. The input file can be a small 100KB file or a massive 4GB file, but either way, you end up with a checksum of the same length. Checksums may also be called hashes. Small changes in the file produce very different looking checksums. You can use checksums to check files and other data for errors that occur during transmission or storage, as well as for evidence in a forensic investigation to ensure it hasn't been tampered with.

Key Topic

Provenance

Data ***provenance*** refers to the establishment of a chain of custody for information that can describe its generation and all subsequent modifications that have led to its current state. An advanced data provenance practice or system gives forensic investigators a transparent idea about the data's lineage and helps resolve disputes over controversial pieces of data by providing digital evidence. Likewise, provenance research is a forensic method employed to reconstruct legal chains of ownership that establish an artwork's whereabouts from the moment of creation to its present circumstances, which is helpful in defending the chain of custody in cases where data is subject to handling scrutiny.

The first step to being able to make use of provenance for forensic purposes is to be able to ensure that it is collected in a secure and trustworthy fashion. However, the

collection process alone raises several significant challenges with approaches to provenance collection from application to operating system level, and this process should rely on a provenance monitor to assure the complete collection of data. Such information can be invaluable for a forensics investigator.

Key Topic Preservation

Handling of digital evidence is the most important aspect in digital forensics. This process is known as ***preservation*** of evidence—isolation and protection of digital evidence exactly as found without alteration so that it can later be analyzed. Sometimes this process involves a forensic copy of a hard disk from a system or logs being collected by another system. Collection is the gathering of devices and duplication of electronically stored information (ESI) for the purpose of preserving digital evidence, an exact copy of the original that remains untouched while digital forensics is performed. In certain cases you may have to turn off the device or isolate ESI in a way that will not alter evidence.

Dead box forensic collection—imaging a device after it is powered off in order to collect digital evidence—still remains an essential part of the digital forensic process. However, *live box forensics* is growing increasingly important with today's technology, with exploit codes capable of remaining in memory. If a device is encrypted, without the passcode or encryption key, you may never have another chance to acquire valuable evidence if that device powers off or locks due to inactivity.

The other problem with live collection is that relevant data could be permanently lost due to continued use of the device. The data must be preserved for collection if it is to be considered for litigation. Time and date stamps will change, system log files will rotate, and valuable information can be lost if you attempt to log in and see what was done before forensic images are taken of a device. A copy of digital evidence must be properly preserved and collected in accordance with forensic best practices. Otherwise, the digital evidence may be inadmissible in court, or spoliation sanctions may be imposed.

Key Topic E-discovery

E-discovery is the process of identifying, preserving, collecting, processing, reviewing, and analyzing electronically stored information in litigation. The digital forensics process involves identifying, preserving, collecting, analyzing, and reporting on digital information. In an e-discovery matter, the role of the expert is to provide the information to legal teams in a reviewable format for the analysis. When leveraging digital forensics, the expert performs the analysis of the information and reports the findings to the legal teams. The party performing the analysis of the electronic information is the primary differentiator between e-discovery and digital forensics.

The terms *e-discovery* and *digital forensics* are often used interchangeably, but there are clear differences. The critical difference is the analysis of the information. In an e-discovery engagement, the legal teams review and analyze the information. In digital forensics, the expert reviews the digital information and provides the findings in an expert report.

TIP Organizations should have a legal hold process in place to perform an e-discovery to preserve and gather information.

Key Topic Data Recovery

Forensic ***data recovery*** is the extraction of data from damaged, deleted, or purposely destroyed evidence sources in a forensically sound manner. This method of recovering data means that any evidence resulting from it can later be relied on in a court of law. Deleting files from your computer doesn't mean the files are necessarily gone, even after being deleted and emptied in the Recycle Bin. What really happens here is that you are deleting the location of those files, which merely hides them from the operating system. Emptying the Recycle Bin frees up disk space and removes the pointer details from the file directory, but the files remain. Unless that data has been physically written over or removed from the hard drive, those files can be recovered with forensic recovery, even if the entire drive has been formatted and seemingly wiped clean.

Key Topic Nonrepudiation

Nonrepudiation is the assurance that someone cannot deny the validity of something, where a statement's author cannot dispute its authorship. Nonrepudiation is a legal concept that is widely used in information security and refers to a service that provides proof of the origin and integrity of the data. Nonrepudiation makes it difficult to successfully deny who and where a message came from as well as the authenticity and integrity of that message. Digital signatures can offer nonrepudiation when it comes to online transactions, where it is crucial to ensure that the party of a contract or a communication can't deny the authenticity of the signature on a document or sending the communication in the first place. In this context, nonrepudiation refers to the ability to ensure that a party to a contract or a communication must accept the authenticity of his or her signature on a document or the sending of a message. In forensics and digital security, nonrepudiation means

A service or system that provides proof of the integrity and origin of data.

An authentication that can be said to be genuine with high confidence.

Proof of data integrity is typically the easiest of these requirements to accomplish. A data hash such as SHA2 usually ensures that the data will not be changed undetectably. Even with this safeguard, it is possible to tamper with data in transit, either through an on-path (formerly known as man-in-the-middle) attack or phishing.

As a result, data integrity is best asserted when the recipient already possesses the necessary verification information, such as after being mutually authenticated. The most common method to provide nonrepudiation in the context of digital communications or storage is through digital signatures, a more powerful tool that provides nonrepudiation in a publicly verifiable manner.

Key Topic

Strategic Intelligence/Counterintelligence

Strategic intelligence is the strategy of narrowing down an investigation to relevant information gathering initially, which is the use of all resources to make determinations that limit the investigation's scope to a manageable level. ***Counterintelligence*** is information gathered and activities conducted to protect against espionage, other intelligence activities, or sabotage conducted by or on behalf of other elements. The intelligence is designed to quickly direct resources to the most significant problems first and address them head on. The strategic intelligence plan should aid in planning and resource allocation, priorities for intelligence protection, and the process.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 30-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 30-2 Key Topics for Chapter 30

Key Topic Element	Description	Page Number
Section	Legal Hold	842
Section	Video	842
Section	Admissibility	843
Section	Chain of Custody	844
Section	Timelines of Sequence of Events	844

Key Topic Element	Description	Page Number
Section	Tags	845
Section	Reports	846
Section	Event Logs	846
Section	Interviews	846
Section	Order of Volatility	848
Section	Disk	848
Section	Random-Access Memory	848
Section	Swap/Pagefile	849
Section	Operating System	850
Section	Device	850
Section	Firmware	851
Section	Snapshot	851
Section	Cache	852
Section	Network	852
Section	Artifacts	853
Section	Right-to-Audit Clauses	854
Section	Regulatory/Jurisdiction	855
Section	Data Breach Notification Laws	855
Section	Hashing	856
Section	Checksums	857
Section	Provenance	857
Section	Preservation	858
Section	E-discovery	858
Section	Data Recovery	859
Section	Nonrepudiation	859
Section	Strategic Intelligence/Counterintelligence	860

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

legal hold, chain of custody, tag, event logs, acquisition, order of volatility, swap file, pagefile, artifacts, right-to-audit, provenance, preservation, E-discovery, data recovery, nonrepudiation, strategic intelligence, counterintelligence

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What are the three rules for evidence?
2. What are three standards for evidence?
3. When are checksums useful?
4. What is the role of a hash in computer forensics?
5. What does NFAT mean in the context of network forensics?

This page intentionally left blank



This chapter covers the following topics related to Objective 5.1 (Compare and contrast various types of controls) of the CompTIA Security+ SY0-601 certification exam:

- Category
  - Managerial
  - Operational
  - Technical
- Control type
  - Preventative
  - Detective
  - Corrective
  - Deterrent
  - Compensating
  - Physical

Comparing and Contrasting the Various Types of Controls

In this chapter you learn about the various categories of controls, starting with an overview of managerial, operational, and technical controls. Additionally, you learn the control types, including preventative, detective, corrective, deterrent, compensating, and physical controls.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 31-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 31-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Control Category	1–3
Control Types	4–7

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which controls are typically addressed by an organization's management?
   a. Managerial controls
   b. Physical controls
   c. Operational controls
   d. None of these answers are correct.

2. Which controls are designed to increase individual and group system security?
   a. Physical controls
   b. Operational controls
   c. Corrective controls
   d. All of these answers are correct.

3. Which controls include auditing and cryptography?
   a. Technical controls
   b. Operational controls
   c. Managerial controls
   d. None of these answers are correct.

4. Which controls are employed before an event and are designed to prevent an incident?
   a. Managerial controls
   b. Compensating controls
   c. Corrective controls
   d. Preventative controls

5. For which of these controls is the primary aim monitoring and detecting any unauthorized behavior?
   a. Corrective controls
   b. Managerial controls
   c. Physical controls
   d. Detective controls

6. Which of the following controls are used after an event?
   a. Physical controls
   b. Corrective controls
   c. Preventative controls
   d. None of these answers are correct.

7. Which of the following are also known as alternative controls?
   a. Physical controls
   b. Compensating controls
   c. Detective controls
   d. None of these answers are correct.

Foundation Topics

Control Category

Controls can be classified into three main categories as described in the sections that follow.

Managerial Controls

Managerial controls are techniques and concerns addressed by an organization's management (managers and executives). Generally, these controls focus on decisions and the management of risk. They also concentrate on procedures, legal and regulatory policies, the software development lifecycle (SDLC), the computer security lifecycle, information assurance, and vulnerability management/scanning. In short, these controls focus on how the security of your data and systems is managed. Managerial or administrative controls include business and organizational processes and procedures, such as security policies and procedures, personnel background checks, security awareness training, and formal change-management procedures.

Operational Controls

Operational controls are the controls executed by people. They are designed to increase individual and group system security. They include user awareness and training, fault tolerance and disaster recovery plans, incident handling, computer support, baseline configuration development, and environmental security. The people who carry out the specific requirements of these controls must have technical expertise and understand how to implement what management desires of them. Operational controls include physical controls that form the outer line of defense against direct access to data, such as protecting backup media; securing output and mobile file storage devices; and paying attention to facility design details, including layout, doors, guards, locks, and surveillance systems.

Technical Controls

Technical controls are the logical controls executed by the computer system. Technical controls include authentication, access control, auditing, and cryptography. The configuration and workings of firewalls, session locks, RADIUS servers, or RAID 5 arrays would be within this category, as well as concepts such as least privilege implementation. Technical/logical controls are security controls put in place that are executed by technical systems. Technical controls include logical access control systems, security systems, encryption, and data classification solutions.

Summary of Control Categories

Table 31-2 provides a quick comparison view of the different control categories.

Key Topic

Table 31-2 Summary of Control Categories

Category	Description
Managerial controls	Managerial or administrative controls include business and organizational processes and procedures, such as security policies and procedures, personnel background checks, security awareness training, and formal change-management procedures.
Operational controls	Operational controls include physical controls that form the outer line of defense against direct access to data, such as protecting backup media; securing output and mobile file storage devices; and paying attention to facility design details, including layout, doors, guards, locks, and surveillance systems.
Technical controls	Technical/logical controls are security controls put in place that are executed by technical systems. Technical controls include logical access control systems, security systems, encryption, and data classification solutions.

Control Types

The following sections focus on control types. They include preventative controls, detective controls, corrective controls, deterrent controls, and compensating controls as well as physical controls.

Preventative Controls

Preventative controls are employed before the event and are designed to prevent an incident. Examples include biometric systems designed to keep unauthorized persons out, network intrusion prevention systems (NIPSs) to prevent malicious activity, and RAID 1 to prevent loss of data. They are also sometimes referred to as deterrent controls. Preventative controls enforce security policy and should prevent incidents from happening. The only way to bypass a preventative control is to find a flaw in its implementation or logic. These controls are usually not optional. Examples of preventative controls are access lists, passwords, and fences. Preventative controls include security awareness, separation of duties, access control, security policies, and intrusion prevention systems.

Detective Controls

Detective controls aim at monitoring and detecting any unauthorized behavior or hazard. These types of controls are generally used to alert a failure in other types

of controls such as preventative, deterrent, and compensating controls. Detective controls are very powerful while an attack is taking place, and they are useful in the post-mortem analysis to understand what has happened. Audit logs, intrusion detection systems (IDSs), motion detection, and Security Information and Event Management (SIEM) systems are examples of detective controls.

These controls are used during an event and can find out whether malicious activity is occurring or has occurred. Examples include CCTV/video surveillance, alarms, network intrusion detection systems (NIDSs), and auditing. Detective controls warn that physical security measures are being violated. Detective controls attempt to identify unwanted events after they have occurred. Common technical detective controls include audit trails, intrusion detection systems, system monitoring, checksums, and anti-malware.

Corrective Controls

Corrective controls are used after an event. They limit the extent of damage and help the company recover from damage quickly. Tape backup, hot sites, and other fault tolerance and disaster recovery methods are also included here. These are sometimes referred to as compensating controls. Corrective controls include all the controls used during an incident to correct the problem. Quarantining an infected computer, sending a guard to block an intruder, and terminating an employee for not having followed the security policy are all examples of corrective controls. Corrective controls are reactive and provide measures to lessen harmful effects or restore the system being impacted. Examples of corrective controls include operating system upgrades, data backup restores, vulnerability mitigation, and antimalware.

Deterrent Controls

Organizations use ***deterrent controls*** to try to deter threat actors from executing an offensive assault on their environment. The idea is that if potential threat actors see that this type of control is in place they may decide to move on. An example of this type of control is an alarm system on a home. If a robber sees that the home clearly has an alarm system, he or she may decide to move on to the next house. In that same sense, threat actors on the Internet may identify that an organization they are targeting is using a specific type of technology that would make it more difficult for them to carry out their attack and may move on.

Deterrent controls are similar to preventative controls in the sense that the primary objective is to prevent an incident from occurring. Unlike preventative controls, however, the rationale behind deterrent controls is to discourage attackers from

proceeding just because a control is in place. For example, a system banner warning that any unauthorized attempt to log in will be monitored and punished is a type of deterrent control. In fact, it would probably discourage casual users from attempting to access the system; however, it might not block determined attackers from trying to log in to the system. Deterrent controls are intended to discourage individuals from intentionally violating information security policies or procedures. And examples of deterrent controls are warnings indicating that systems are being monitored.

Compensating Controls

Also known as alternative controls, ***compensating controls*** are mechanisms put in place to satisfy security requirements that are either impractical or too difficult to implement. For example, instead of using expensive hardware-based encryption modules, an organization might opt to use network access control (NAC), data loss prevention (DLP), and other security methods. Or, on the personnel side, instead of implementing separation of duties, an organization might opt to do additional logging and auditing. You should approach compensating controls with great caution. They do not give the same level of security as their replaced counterparts.

Of course, many security concepts can be placed in the category of physical as well as other categories listed in the sections that follow. For example, a locking door would be an example of a physical control as well as a preventative control.

When you see technologies, policies, and procedures in the future, attempt to place them within their proper control category. Semantics will vary from one organization to the next, but as long as you can categorize security features in a general fashion such as the ones previously listed, you should be able to define and understand just about any organization's security controls. Compensating controls are alternate controls that are intended to reduce the risk of an existing or potential control weakness. They can include audit trails and transaction logs that someone in a higher position reviews.

Physical Controls

To control access, ***physical controls*** can be considered the first line of defense, sort of like a firewall is the first line of defense for a network. Implementing physical access security methods should be a top priority for an organization. Unfortunately, securing physical access to the organization's building sometimes slumps to the bottom of the list. Or a system is employed, but it fails to mitigate risk properly. In some cases, the system is not maintained well. Proper building entrance access and secure access to physical equipment are vital. And anyone coming and going should be logged and surveyed.

Operational/physical controls include organizational culture and physical controls that form the outer line of defense against direct access to data, such as protecting backup media; securing output and mobile file storage devices; and paying attention to facility design details, including layout, doors, guards, locks, and surveillance systems. For more in-depth information, physical security is covered in detail in Chapter 15, "Understanding the Importance of Physical Security Controls."

Summary of Control Types

Table 31-3 provides a quick comparison view of the different control categories.

Key Topic

Table 31-3 Summary of Control Types

Control type	Description
Preventative	Preventative controls include security awareness, separation of duties, access control, security policies, and intrusion prevention systems.
Detective	Detective controls warn that physical security measures are being violated. Detective controls attempt to identify unwanted events after they have occurred. Common technical detective controls include audit trails, intrusion detection systems, system monitoring, checksums, and antimalware.
Corrective	Corrective controls are reactive and provide measures to lessen harmful effects or restore the system being impacted. Examples of corrective controls include operating system upgrades, data backup restores, vulnerability mitigation, and antimalware.
Deterrent	Deterrent controls are intended to discourage individuals from intentionally violating information security policies or procedures. Examples of deterrent controls are warnings indicating that systems are being monitored.
Compensating	Compensating controls are alternate controls that are intended to reduce the risk of an existing or potential control weakness. These can include audit trails and transaction logs that someone in a higher position reviews.
Physical	Operational/physical controls include organizational culture and physical controls that form the outer line of defense against direct access to data, such as protecting backup media; securing output and mobile file storage devices; and paying attention to facility design details, including layout, doors, guards, locks, and surveillance systems.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 31-4 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 31-4 Key Topics for Chapter 31

Key Topic Element	Description	Page Number
Table 31-2	Summary of Control Categories	869
Table 31-3	Summary of Control Types	872

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

managerial controls, operational controls, technical controls, preventative controls, detective controls, corrective controls, deterrent control, compensating controls, physical controls

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What control category is addressed by an organization's management?
2. What control category is designed to increase individual and group system security?
3. What control category would include firewalls?
4. What control type enforces security policy?
5. What control type is intended to discourage someone from violating policies?
6. What control type warns that physical security measures are being violated?
7. What control type includes all the controls used during an incident?
8. What control type is also known as an alternative control?
9. Which type of control would include something like door access?
10. Which type of control would you put in place to control access to a server room?



This chapter covers the following topics related to Objective 5.2 (Explain the importance of applicable regulations, standards, or frameworks that impact organizational security posture) of the CompTIA Security+ SY0-601 certification exam:

- Regulations, standards, and legislation
  - General Data Protection Regulation (GDPR)
  - National, territory, or state laws
  - Payment Card Industry Data Security Standard (PCI DSS)
- Key frameworks
  - Center for Internet Security (CIS)
  - National Institute of Standards and Technology (NIST) Risk Management Framework (RMF)/ Cybersecurity Framework (CSF)
  - International Organization for Standardization (ISO) 27001/27002/27701/31000
  - SSAE SOC 2 Type I/II
  - Cloud security alliance
  - Cloud controls matrix
  - Reference architecture
- Benchmarks /secure configuration guides
  - Platform/vendor-specific guides
    - Web server
    - OS
    - Application server
    - Network infrastructure devices

CHAPTER 32

Understanding the Importance of Applicable Regulations, Standards, or Frameworks That Impact Organizational Security Posture

In this chapter you learn how to understand the importance of applicable regulations, standards, or frameworks that impact organizational security posture. The chapter starts by examining regulations, standards, and legislation. This discussion includes the General Data Protection Regulation (GDPR) as well as national, territory, or state laws. The chapter also covers the Payment Card Industry Data Security Standard (PCI DSS) and briefly digs into some key frameworks, including the Center for Internet Security (CIS), National Institute of Standards and Technology (NIST) RMF/CSF, and the International Organization for Standardization (ISO) 27001/27002/27701/31000. The chapter also touches on the SSAE SOC 2 Type I/II and Cloud Security Alliance including the cloud security matrix and reference architecture. The chapter concludes with a discussion of benchmarks and security configuration guides, such as platform- and vendor-specific guides for web servers, operating systems, application servers, and network infrastructure devices.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 32-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 32-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Regulations, Standards, and Legislation	1–3
Key Frameworks	4–7
Benchmarks and Secure Configuration Guides	8

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which regulation was established in the European Union?
   a. GDPR
   b. SOX
   c. GLB
   d. None of these answers are correct.

2. Which act governs the disclosure of financial and accounting information?
   a. GLB
   b. SOX
   c. HIPAA
   d. All of these answers are correct.

3. Which act governs the disclosure and protection of health information?
   a. HIPAA
   b. GLB
   c. SOX
   d. None of these answers are correct.

4. Which of the following is a nonprofit organization established in 2000 to provide security best practice guidance for enhancing the security of cyberspace?
   a. Cloud Security Alliance
   b. International Organization for Standardization
   c. National Institute of Standards and Technology
   d. Center for Internet Security

5. Which of the following developed the Risk Management Framework in 2017?
   a. Cloud Security Alliance
   b. International Organization for Standardization
   c. National Institute of Standards and Technology
   d. Center for Internet Security

6. Which of the following developed the Cybersecurity Framework in 2014?
   a. Cloud Security Alliance
   b. International Organization for Standardization
   c. National Institute of Standards and Technology
   d. Center for Internet Security

7. Which of the following is a nonprofit organization established in 2008 with the goal of promoting security best practices in cloud computing environments?
   a. Cloud Security Alliance
   b. International Organization for Standardization
   c. National Institute of Standards and Technology
   d. Center for Internet Security

8. Which of the following was created to provide a standardized solution for security automation?
   a. SOX
   b. GLB
   c. SCAP
   d. OVAL

Foundation Topics

Regulations, Standards, and Legislation

There are myriad legislative laws and policies. For the Security+ exam, we are concerned only with a few that affect—and protect—the privacy of individuals. Here, we cover those and some associated security standards.

More important for the Security+ exam are organizational policies. Organizations usually define policies that concern how data is classified, how employees are expected to behave, and how to dispose of IT equipment that is no longer needed. These policies begin with a statement or goal that is usually short, to the point, and open-ended. These statements are normally written in clear language that can be understood by most everyone. They are followed by procedures (or guidelines) that detail how the policy will be implemented.

Table 32-2 shows an example of a basic policy and corresponding procedure.

Table 32-2 Example of a Company Policy

Policy	Procedure
Employees will identify themselves in a minimum of two ways when entering the complex.	1. When employees enter the complex, they will first enter a guard room. This will begin the authentication process. 2. In the guard room, they must prove their identification in two ways: ■ By showing their ID badge to the on-duty guard. ■ By being visible to the guard so that the guard can compare their likeness to the ID badge's photo. The head of the employee should not be obstructed by hats, sunglasses, and so on. In essence, the employee should look similar to the ID photo. If the employee's appearance changes for any reason, that person should contact human resources for a new ID badge. If guards cannot identify the "employee," they will contact the employee's supervisor, human resources, or security in an attempt to confirm the person's identity. If the employee is not confirmed, that person will be escorted out of the building by security. 3. After the guard has acknowledged the identification, employees will swipe their ID badge against the door scanner to complete the authentication process and gain access to the complex.

Keep in mind that this is just a basic example; technical documentation specialists will tailor the wording to fit the feel of the specific organization. Plus, the procedure will be different depending on the size and resources of the organization and the type of authentication scheme used, which could be more or less complex. However, the *policy* (which is fairly common) is written in such a way as to be open-ended, allowing for the *procedure* to change over time. We discuss many different policies as they relate to the Security+ exam in the following sections. A regulation that is top of mind for many organizations these days is the General Data Protection Regulation.

Key Topic

General Data Protection Regulation

The ***General Data Protection Regulation (GDPR)*** is a European Union (EU) law that was enacted in 2018. Its overall focus is on data protection and privacy for individuals. Although it is a law enacted in the EU, it applies to any organization collecting information about people in the EU. That means if your organization collects and handles the personal data of EU citizens, then this regulation applies to you. For instance, if you run a business that offers goods or services in the EU and that business requires you to collect information about your customers, then you would be required to follow the GDPR requirements. Not following them could result in large fines. This is one of the factors that makes GDPR a larger concern to organizations than many other laws that have been in place for many years. GDPR penalties can be very high. They are no slap on the wrist. For additional information on GDPR, refer to https://gdpr.eu/.

Key Topic

National, Territory, or State Laws

A number of laws are enforced at different levels of government and for different geographical regions. A majority of them are enacted at the federal level. Table 32-3 describes some of the various laws and regulations that have been passed concerning the disclosure of data.

Table 32-3 Acts Passed Concerning the Disclosure of Data, Personally Identifiable Information (PII), and Protected Health Information (PHI)

Act	Acronym	Description
Privacy Act of 1974	n/a	Establishes a code of fair information practice. Governs the collection, use, and dissemination of personally identifiable information about persons' records maintained by federal agencies.

Act	Acronym	Description
Sarbanes–Oxley	SOX	Governs the disclosure of financial and accounting information. Enacted in 2002.
Health Insurance Portability and Accountability Act	HIPAA	Governs the disclosure and protection of health information (PHI). Enacted in 1996.
Gramm-Leach-Bliley Act	GLB	Enables commercial banks, investment banks, securities firms, and insurance companies to consolidate. Enacted in 1999. Protects against pretexting. Individuals need proper authority to gain access to nonpublic information such as Social Security numbers.
Help America Vote Act of 2002	HAVA	Replaced punch card and lever-based voting systems. Governs the security, confidentiality, and integrity of personal information collected, stored, or otherwise used by various electronic and computer-based voting systems.
California SB 1386	SB 1386	Requires California businesses that store computerized personal information to immediately disclose breaches of security. Enacted in 2003.
Personal Information Protection and Electronic Documents Act	PIPEDA	Protects personal information of Canadian citizens. It applies to any private-sector organization in Canada that collects, uses, or discloses personal information.

Many computer technicians have to deal with SOX and HIPAA at some point in their careers, and although these types of acts create a lot of paperwork and protocols, the expected result is that, in the long run, they will help companies protect their data and keep sensitive information private.

NOTE SOX sparked another concept known as governance, risk, and compliance (GRC), which deals with the continuous security monitoring and the overall management of information systems and control structures; risk management processes; and compliance with stated requirements, whether they are government related or otherwise.

Payment Card Industry Data Security Standard (PCI DSS)

Key Topic

The ***Payment Card Industry Data Security Standard (PCI DSS)*** was enacted in 2006 as a joint effort by credit card industry partners such as Visa, Mastercard, American Express, Discover, and JCB. The overall goal of PCI DSS is to enhance the security around payment card data processing. The requirement applies to any organization that processes payment card data and enforces penalties for noncompliance on such organizations.

PCI compliance requirements can be summed up as the following:

- Protect cardholder data.
- Build and maintain a secure network.
- Maintain an information security policy.
- Maintain a vulnerability management program.
- Implement strong access control measures.
- Regularly monitor and test systems and networks.

Key Frameworks

Here we discuss some key frameworks, beginning with an overview of the Center for Internet Security (CIS). From there we discuss the National Institute of Standards and Technology (NIST) RMF/CSF. The discussion continues with a look at the International Organization for Standardization (ISO) 27001/27002/27701/31000 as well as SSAE SOC 2 Type I/II. We conclude with a discussion of the Cloud Security Alliance including the Cloud Controls Matrix and reference architecture.

Although an organization might opt to create its own framework, it makes sense for organizations—especially larger ones—to use standards that have already been thoroughly planned out, or at least to base their framework on those standards; for example, the ISO/IEC 27000 family of information security standards. In addition to ISO/IEC 27002:2013, you will find several others. You can find more information at www.iso.org/isoiec-27001-information-security.html.

Then, of course, there is NIST, which defines all kinds of guidelines and recommendations within the SP 800 and SP 1800 publication groups. For more information, see http://csrc.nist.gov/publications/PubsSPs.html.

NOTE NIST introduced the NIST Cybersecurity Framework (CSF), which resulted from a collaboration between the government and the private sector. CSF organizes cybersecurity activities at their highest level. The highest levels are called functions, including Identify, Protect, Detect, Respond, and Recovery. Control Objectives for Information and Related Technology (COBIT) is a set of best practices for IT management. COBIT is most commonly used to attain compliance with the Sarbanes–Oxley Act (SOX). ISO/IEC 27001 is a standard for information security management, for which organizations may be certified if they meet the requirements. ISO/IEC 27701 extends ISO 27001 with enhancements for privacy in order to establish and maintain information management systems specific to privacy. ISO/IEC 31000 provides a framework for the risk management process.

Next, there is ISACA's COBIT framework, which divides IT into four sections: (1) plan and organize; (2) acquire and implement; (3) deliver and support; and (4) monitor and evaluate. That pretty much sums up everything we've described in this book!

Also, you might be interested in the Information Technology Infrastructure Library (ITIL), Business Information Services Library (BiSL), and Project Management Body of Knowledge (PMBOK). You can find a good NIST document that combines the use of several of these at www.nist.gov/system/files/documents/cyberframework/cybersecurity-framework-021214.pdf. You will also find that the U.S. government and military have their own resources on the subject, or depending on the scenario, will use one of the aforementioned standards.

Some of these frameworks are regulatory, and you as an employee *must* abide by any of them that are applicable to your organization or profession. Some are nonregulatory, but usually the organization strongly urges its employees to accept them.

Reference frameworks can also be industry-specific or could define how precise tasks and problems within an organization are to be approached. For example, your company might repair mobile devices for corporations. This company would require a specific secure configuration guide detailing how the mobile devices are repaired, stored, handled, and so on. Or, you might be interested in benchmarking your servers. A detailed list of procedures is vital so that you obtain reliable results in a controlled environment. Then there is software development: When building software, you might embrace the concept of *use case analysis*, which is a requirement analysis technique practiced in software engineering. The use case analysis can benefit from well-written procedures within an IT security framework. Let's not forget about software-defined networking (SDN), which is an approach to computer networking that allows administrators to programmatically control and manage network behavior via open interfaces such as OpenFlow and Cisco's Open Network Environment. SDN can benefit greatly from a well-thought-out framework.

Your IT security framework might include risk analysis and vulnerability assessment tools and how to use them—for example, using the Security Content Automation Protocol (SCAP) to automate vulnerability management. The framework might also incorporate how to properly utilize enterprise resource planning (ERP) software, which is used to manage and automate many back-office technology functions in larger organizations. The examples are endless; just about anything we've described in this book can be incorporated into your IT security framework.

Consequently, the IT security framework could be large or small. It might deal with a specific task or many tasks within an organization. But often, the content in the framework can be applied to many different solutions and implementations. The goal is to organize a group of processes, procedures, and policies from your organization into a single cohesive agenda that all employees can easily understand and work within.

From a security perspective, the IT security framework—if designed properly—can help the organization provide for defense in depth of systems and networks, and increase the confidentiality, integrity, and availability of data.

The Statement on Standards for Attestation Engagements (SSAE) is a standard for the auditing of service organizations. It is produced by the American Institute of Certified Public Accountants (AICPA). The latest version as of this writing is SSAE 18, which superseded SSAE 16. You might be familiar with the Statement on Auditing Standards, also known as SAS 70, which has been around for a while. Many organizations have moved from utilizing SAS 70 to SSAE. SSAE 18 includes two different types of reports. An SOC 1 report is related to the secure handling of financial information about the company. SOC 1 has two different types of reports. A Type 1 report is meant to show that your company's financial controls are designed properly, and a Type 2 report is meant to show that those controls are able to work effectively over time. SOC 2, which is more relevant to this book and further discussions, is a framework that an organization uses to demonstrate that its cloud and data center security controls are sufficient. This is something that you should expect from any service provider that you are doing business with.

Table 32-4 outlines the key IT security frameworks.

Key Topic

Table 32-4 Key Frameworks

Framework	Acronym	Description
Center for Internet Security	CIS	CIS is a nonprofit organization that was established in 2000. Its overall goal is to provide security best practice guidance for enhancing the security of cyberspace.

Framework	Acronym	Description
National Institute of Standards and Technology	(NIST) RMF/CSF	NIST developed the Risk Management Framework (RMF) in 2017 as a result of an executive order from the president, which required all federal agencies to comply with it. For more information about the NIST RMF, visit https://www.nist.gov/cyberframework/risk-management-framework. NIST developed the Cybersecurity Framework (CSF) in 2014 as a result of executive order 13636 from the president. The NIST CSF is made of up four core framework functions: Identify, Protect, Detect, and Respond and Recover. For more information about the NIST CSF, visit www.nist.gov/cyberframework.
International Organization for Standardization	(ISO) 27001/27002/27701/31000	ISO was originally established in 2005 and later updated in 2013. It addresses the following topics: ■ The organizational context ■ Involvement of the leadership ■ Planning and objectives ■ Support, including resources and communication ■ Operational aspects ■ Evaluation of performance ■ Continuous improvement
Standards of Attestations Engagement (SSAE) System and Organization Controls (SOC)	SSAE SOC 2 Type I/II	The main focus of the Standards of Attestations Engagement (SSAE) System and Organization Controls (SOC) is related to the controls internal to an organization that are related to security compliance and operations. It is built on the concept of trust principles. For more information related to SSAE SOC 2, visit www.ssae-16.com/soc-1-report/the-ssae-18-audit-standard/.
Cloud Security Alliance	CSA	CSA is a nonprofit organization established in 2008 with the goal of promoting security best practices in cloud computing environments. The ***Cloud Controls Matrix*** is a framework established by the CSA for cloud computing. The organization also develops the ***reference architecture*** to help cloud providers with guidance on developing secure interoperability best practices. For more information related to CSA, visit https://cloudsecurityalliance.org/.

Benchmarks and Secure Configuration Guides

In the following sections, we discuss benchmarks and secure configuration guides, including platform- and vendor-specific guides such as web server, operating system, application server, and network infrastructure devices. These guides are useful when hardening your systems and network. They can also be utilized to obtain compliance for regulatory purposes.

Security Content Automation Protocol

Security Content Automation Protocol (SCAP) was created to provide a standardized solution for security automation. The SCAP mission is to maintain system security by ensuring security configuration best practices are implemented in the enterprise network, verifying the presence of patches, and maintaining complete visibility of the security posture of systems and the organization at all times.

The current SCAP specifications include the following:

- **Languages:**
  - **Open Vulnerability and Assessment Language (OVAL):** OVAL is an international community standard to promote open and publicly available security content and to standardize the transfer of this information in security tools and services. More information about OVAL is available at https://oval.mitre.org.
  - **Extensible Configuration Checklist Description Format (XCCDF):** XCCDF is a specification for a structured collection of security checklists and benchmarks. More information about XCCDF is available at https://scap.nist.gov/specifications/xccdf.
  - **Open Checklist Interactive Language (OCIL):** OCIL is a framework for collecting and interpreting responses from questions offered to users. More information about OCIL is available at https://scap.nist.gov/specifications/ocil.
  - **Asset Identification (AI):** AI is a specification designed to quickly correlate different sets of information about enterprise computing assets. More information about AI is available at https://scap.nist.gov/specifications/ai.
  - **Asset Reporting Format (ARF):** ARF is a specification that defines the transport format of information about enterprise assets and provides a standardized data model to streamline the reporting of such information. More information about ARF is available at https://scap.nist.gov/specifications/arf.

- **Enumerations:**
  - **Common Vulnerabilities and Exposures (CVE):** CVE assigns identifiers to publicly known system vulnerabilities. Cisco assigns CVE identifiers to security vulnerabilities according to the Cisco public vulnerability policy at https://www.cisco.com/web/about/security/psirt/security_vulnerability_policy.html. More information about CVE is available at https://cve.mitre.org.
  - **Common Platform Enumeration (CPE):** CPE is a standardized method of naming and identifying classes of applications, operating systems, and hardware devices. More information about CPE is available at https://nvd.nist.gov/cpe.cfm.
  - **Common Configuration Enumeration (CCE):** CCE provides unique identifiers for configuration guidance documents and best practices. The main goal of CCE is to enable organizations to perform fast and accurate correlation of configuration issues in enterprise systems. More information about CCE is available at https://nvd.nist.gov/cce/index.cfm.

NOTE Other community-developed enumerators, such as the Common Weakness Enumeration (CWE), are currently being expanded and further developed. CWE is a dictionary of common software architecture, design, code, or implementation weaknesses that could lead to security vulnerabilities. More information about CWE is available from http://cwe.mitre.org. Another emerging enumerator is the Common Remediation Enumeration (CRE). More information about CRE is available at http://scap.nist.gov/specifications/cre.

- **Metrics:**
  - **Common Vulnerability Scoring System (CVSS):** CVSS is a standards-based scoring method that conveys vulnerability severity and helps determine the urgency and priority of response. You can obtain the latest CVSS specification documentation, examples of scored vulnerabilities, and a calculator at first.org/cvss.
  - **Common Configuration Scoring System (CCSS):** More information about CCSS is available in the following PDF document: https://csrc.nist.gov/publications/nistir/ir7502/nistir-7502_CCSS.pdf.

NOTE Two emerging metrics specifications are the *Common Weakness Scoring System* (CWSS) and the *Common Misuse Scoring System* (CMSS). CWSS is a methodology for scoring software weaknesses. CWSS is part of CWE. More information about CWSS is available at https://cwe.mitre.org/cwss. CMSS is a standardized way to measure software feature misuse vulnerabilities. More information about CMSS is available at https://scap.nist.gov/emerging-specs/listing.html#cmss.

- **Integrity:** Integrity is provided by the *Trust Model for Security Automation Data* (TMSAD), which is a trust model for maintaining integrity, authentication, and traceability of security automation data. More information about TMSAD is available in the following PDF document: https://csrc.nist.gov/publications/nistir/ir7802/NISTIR-7802.pdf.

Figure 32-1 summarizes the SCAP components.

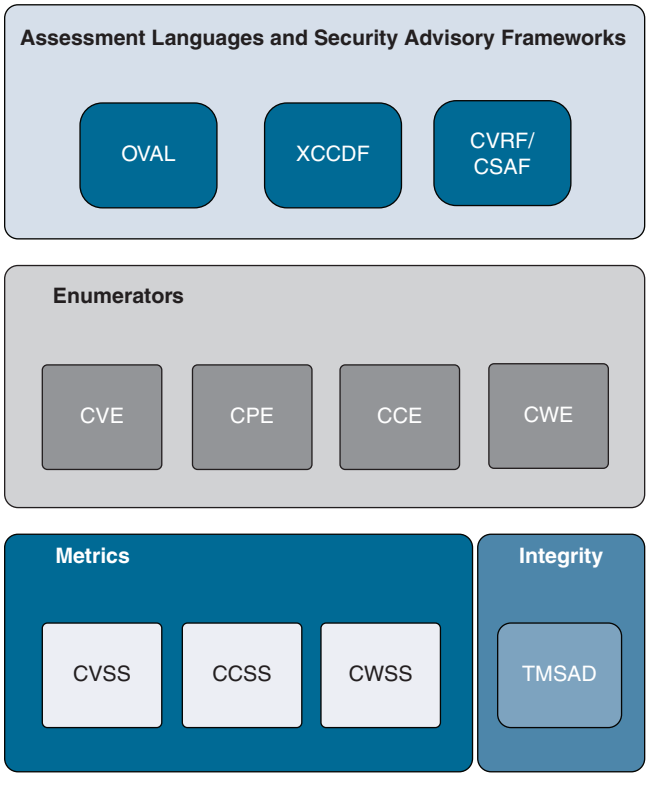


FIGURE 32-1 SCAP Components

Table 32-5 summarizes the benchmarks and secure configuration guides for platforms/vendors, web servers, operating systems, application servers, and network infrastructure devices.

Key Topic

Table 32-5 Benchmarks/Secure Configuration Guides

Framework	Description	Examples of Guides Available
Platform/vendor-specific guides	Some guides are specific to vendor products, and many times are documented and provided by the vendors themselves.	Cisco Configuration guides available on Cisco.com: www.cisco.com/c/en/us/support/all-products.html Red Hat Enterprise Linux Security Guide: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/index
Web server	Web server security hardening guides are available from organizations like the Center for Internet Security and DISA Security Implementation Guidelines.	Center for Internet Security (CIS) Apache Benchmark: www.cisecurity.org/benchmark/apache_http_server/ DISA Apache Server Security Technical Implementation Guide: www.stigviewer.com/stig/apache_server_2.4_unix_site/
Operating systems	Operating system security hardening guides are available from organizations like the Center for Internet Security and DISA Security Implementation Guidelines, as well as specific OS vendors.	Center for Internet Security (CIS) Windows Server Benchmark: www.cisecurity.org/benchmark/microsoft_windows_server/ DISA Technical Implementation Guides: https://public.cyber.mil/stigs/
Application server	Application server security hardening guides are available from organizations like the Center for Internet Security and DISA Security Implementation Guidelines.	Oracle Hardening Guide: https://docs.oracle.com/cd/E25178_01/fusionapps.1111/e16690/F371476AN1062D.htm DISA Technical Implementation Guides: https://public.cyber.mil/stigs/
Network infrastructure devices	Network infrastructure device hardening guides are available from organizations like the Center for Internet Security and DISA Security Implementation Guidelines, as well as specific device vendors.	VMWare Security Hardening Guides: www.vmware.com/security/hardening-guides.html Cisco CIS Benchmark: www.cisecurity.org/benchmark/cisco/ Cisco Security Hardening Guide: www.cisco.com/c/en/us/support/docs/ip/access-lists/13608-21.html

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 32-6 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 32-6 Key Topics for Chapter 32

Key Topic Element	Description	Page Number
Section	General Data Protection Regulation (GDPR)	879
Section	National, Territory, or State Laws	879
Paragraph	Description of Payment Card Industry Data Security Standard (PCI DSS)	881
Table 32-4	Key Frameworks	883
Table 32-5	Benchmarks/Secure Configuration Guides	888

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

General Data Protection Regulation (GDPR), Privacy Act of 1974, Sarbanes–Oxley, Health Insurance Portability and Accountability Act, Gramm-Leach-Bliley Act, Help America Vote Act of 2002, California SB 1386, Payment Card Industry Data Security Standard (PCI DSS), Center for Internet Security, National Institute of Standards and Technology, International Organization for Standardization, SSAE SOC 2 Type I/II, Cloud Security Alliance, Cloud Controls Matrix, reference architecture

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What regulation was established in the European Union to protect data and privacy?
2. What act governs the disclosure of financial and accounting information?
3. What act governs the disclosure and protection of health information?
4. Which nonprofit organization enstablished in 2000 focuses on security best practices guides?
5. Which organization developed the Cybersecurity Framework in 2014?
6. Which nonprofit organization established in 2008 is focused on cloud security best practices?
7. _____ was created to provide a standardized solution for security automation.

This page intentionally left blank



This chapter covers the following topics related to Objective 5.3 (Explain the importance of policies to organizational security) of the CompTIA Security+ SY0-601 certification exam:

- Personnel
  - Acceptable use policy
  - Job rotation
  - Mandatory vacation
  - Separation of duties
  - Least privilege
  - Clean desk space
  - Background checks
  - Non-disclosure agreement (NDA)
  - Social media analysis
  - Onboarding
  - Offboarding
  - User training
    - Gamification
    - Capture the flag
    - Phishing campaigns
    - Phishing simulations
    - Computer-based training (CBT)
    - Role-based training
- Diversity of training techniques
- Third-party risk management
  - Vendors
  - Supply chain
  - Business partners
  - Service-level agreement (SLA)
  - Memorandum of understanding (MOU)
  - Measurement systems analysis (MSA)
  - Business partnership agreement (BPA)
  - End of life (EOL)
  - End of service life (EOSL)
  - NDA
- Data
  - Classification
  - Governance
  - Retention
- Credential policies
  - Personnel
  - Third-party
  - Devices
  - Service accounts
  - Administrator/root accounts
- Organizational policies
  - Change management
  - Change control
  - Asset management

CHAPTER 33

Understanding the Importance of Policies to Organizational Security

Some smaller companies don't have much in the way of policies. Arguably, that is why a percentage of them fail. You will see many companies of all sizes create their own policies or embrace ones that other organizations are using, or perhaps apply for a company-wide standards certification from an organization such as the International Organization for Standardization (ISO). Some organizations are bound by legislative policy and organized protocols. In general, policies are designed to protect employees and make the organization more productive and efficient.

Before we continue, it's important to distinguish between policies and procedures. A *policy* is something that an individual employee or entire organization should adhere to but is usually expressed in broad terms. A *procedure* is usually much more specific. Although it is often stated in detail, it can potentially be interpreted more loosely. Standard operating procedures used by corporations, government, and the military are usually pretty tight. But other procedures might be a bit more relaxed, and as long as employees get to the final goal efficiently, procedures can often be overlooked to a certain degree. However, incident response procedures—once developed by an organization—are usually followed to the letter. Otherwise, there can be legal repercussions. Keep in mind that a procedure could be a part of an overall policy.

To help organize the many procedures and policies, you need a plan. An IT security framework is just that: it's like the blueprint for your organization's security goals. It defines, organizes, and interconnects the various policies and procedures that can make people giddy.

The concepts in this chapter are meant to oversee everything else in the book from a more managed perspective. By using a well-planned IT security framework (or frameworks), your procedures and policies—and technology in general—all start to flow together.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 33-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 33-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Personnel Policies	1–6
Diversity of Training Techniques	7
Third-Party Risk Management	8–9
Data Concepts	10
Credential Policies	11
Organizational Policies	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. In which types of organizational policies would you find information on handling of personal information?
   a. Privacy policies
   b. Public policies
   c. Wireless policies
   d. None of these answers are correct.

2. Which type of policy would define rules that restrict how a computer may be used?
   a. Privacy policy
   b. Acceptable use policy
   c. Credential policy
   d. All of these answers are correct.

3. Which of the following describes when more than one person is required to complete a particular task or operation?
   a. Separation of duties
   b. Acceptable use
   c. Mandatory vacation
   d. Job rotation

4. Which concept includes ensuring that IT infrastructure risks are known and managed?
   a. Separation of duties
   b. Job rotation
   c. Offboarding
   d. Due diligence

5. _____ is the mitigation action that an organization takes to defend against the risks that have been uncovered during due diligence.
   a. Acceptable use
   b. Separation of duties
   c. Due process
   d. Due care

6. Which of the following is the principle that an organization must respect to safeguard personnel's rights?
   a. Clean desk space
   b. Due process
   c. Least privilege
   d. User training

7. Which of the following is information used to uniquely identify, contact, or locate a person?
   a. PII
   b. EOL
   c. EOSL
   d. CBT

8. Which of the following is a letter of intent, in regard to a BPA?
   a. Interconnection security agreement (ISA)
   b. Business partnership agreement (BPA)
   c. Memorandum of understanding (MOU)
   d. None of these answers are correct.

9. Which type of contract can establish the profits each partner will get?
   a. Interconnection security agreement (ISA)
   b. Business partnership agreement (BPA)
   c. Memorandum of understanding (MOU)
   d. None of these answers are correct.

10. Which classification is typically used in government and represents the highest sensitivity of data?
    a. Confidential
    b. Top secret
    c. Secret
    d. None of these answers are correct.

11. Which type of account is not tied to a specific user but would be critical for a web server to run?
    a. User
    b. Service
    c. Computer
    d. None of these answers are correct.

12. Which is a structured way of changing the state of a computer system, network, or IT procedure?
    a. Asset management
    b. Change management
    c. Policy management
    d. None of these answers are correct.

Foundation Topics

Personnel Policies

Most organizations have policies governing employees. The breadth and scope of these policies vary from organization to organization. For example, a small company might have a few pages defining how employees should behave (a code of ethics) and what to do in an emergency. Larger organizations might go so far as to certify to a particular standard such as ISO 9001:2015. This means that the organization will comply with a set of quality standards that is all-encompassing, covering all facets of the business. An organization would have to be examined and finally accredited by an accrediting certification body to state that it is ISO certified. This rigorous process is not for the average organization. For many companies, this process would create too much documentation and would bog the company down in details and minutia.

As an IT person, you are more interested in policies that deal with the security of the infrastructure and its employees. As a security administrator, you might deal with procedural documentation specialists, technical documentation specialists, and even outside consultants. You should become familiar with policies and as many procedures as possible, focusing on policies that take security into account, but remember that actual work must take precedence!

Let's look at a few types of policies that are common to organizations. We focus on the security aspect of these policies.

Privacy Policies

The Privacy Act of 1974 sets many standards when it comes to the security of personally identifiable information (PII). However, most organizations will go further and define their own privacy policy, which explains how users' identities and other similar information will be secured. For example, if an organization has an Internet-based application that internal and external users access, the application will probably retain some of their information—possibly details of their identity. Not only should this information be secured, but the privacy policy should state in clear terms what data is allowed to be accessed, and by whom, as well as how the data will be retained and distributed (if at all). An organization might also enact a policy that governs the labeling of data to ensure that all employees understand what data they are handling and to prevent the mishandling of confidential information. Before any system administrators or other personnel gather information about these users, they should consult the privacy policy.

Acceptable Use

Acceptable use policies (AUPs) define the rules that restrict how a computer, network, or other system may be used. They state what users are and are not allowed to do when it comes to the technology infrastructure of an organization. Often, employees must sign an AUP before they begin working on any systems. This policy protects the organization but also defines to employees exactly what they should and should not be working on. If a director asks a particular employee to repair a particular system that was outside the AUP parameters, the employee would know to refuse. If employees are found working on a system that is outside the scope of their work, and they signed an AUP, this is grounds for termination. As part of an AUP, employees enter into an agreement acknowledging they understand that the unauthorized sharing of data is prohibited. Also, employees should understand that they are not to take any information or equipment home without express permission from the various parties listed in the policy. This policy can sometimes be in conflict with a bring-your-own-device (BYOD) policy where users are permitted to bring their own devices into work and use them for work purposes. At that point, strong policies for data ownership need to be developed, identifying what portion of the data on a mobile device is owned by the organization, and what portion is owned by the employee. Any organizational data on a mobile device should be backed up.

Separation of Duties/Job Rotation

Separation of duties defines when more than one person is required to complete a particular task or operation. This policy distributes control over a system, infrastructure, or particular task. ***Job rotation*** is one of the checks and balances that might be employed to enforce the proper separation of duties. It is when two or more employees switch roles at regular intervals. It is used to increase user insight and skill level, and to decrease the risk of fraud and other illegal activities. Both of these policies are enforced to increase the security of an organization by limiting the amount of control a person has over a situation and by increasing employees' knowledge of what other employees are doing.

Mandatory Vacations

Some organizations require employees to take *X* number of consecutive days of vacation over the course of a year as part of their annual leave. For example, a company might require an IT director to take five consecutive days' vacation at least once per year to force another person into that role for that time period. Although a company might state that this policy helps the person rest and focus on the job, and incorporate job rotation, the underlying security concept is that ***mandatory***

vacations can help stop any possible malicious activity that might occur such as fraud, sabotage, or embezzlement. Because IT people are smart and often access the network remotely in a somewhat unobserved fashion, auditing becomes very important.

Onboarding and Offboarding

Onboarding is the process of adding a new employee to an organization and to its identity and access management system. This process incorporates ***user training***, formal meetings, lectures, and human resources employee handbooks and videos. Depending on the organization, it may also include ***background checks*** and ***social media analysis***. It can also be implemented when a person changes roles within an organization. It is known as a socialization technique used to ultimately provide better job performance and higher job satisfaction. Onboarding is associated with federated identity management. It is also sometimes connected to an employee's *role* in the company and therefore role-based access control (RBAC).

NOTE Many organizations use social media analysis to research a potential employee or volunteer. The procedure to do this includes investigating the person's online presence, including Facebook, Instagram, and LinkedIn.

Offboarding is the converse of onboarding; it correlates to procedurally removing an employee from a federated identity management system, restricting rights and permissions, and possibly debriefing the person or conducting an exit interview. This step is taken when a person changes roles within an organization or departs the organization altogether.

An organization will commonly work with business partners, but no business relationship lasts forever, and new ones are often developed. So, onboarding and offboarding also can apply to business partners. The main concern is access to data. Extranets are commonly used technologies with business partners. These technologies allow an organization to carefully select which data the business partner has access to. As relationships with business partners are severed, a systematic audit of all shared data should be made, including the various types of connectivity, permissions, policies, and even physical access to data.

A clean desk policy (CDP) directs employees as to how they should leave their working environment when they are not physically there to monitor it. For instance, they should leave the desk space free of any paperwork that might contain sensitive information. This policy might also involve logging off any computing devices or phones.

Personnel Security Policies

Table 33-2 breaks down the various personnel security policy types.

Key Topic

Table 33-2 Policy Types

Type	Description
Acceptable use policy (AUP)	This policy defines the rules that restrict how a computer, network, or other system may be used.
Separation of duties	This policy specifies when more than one person is required to complete a task.
Job rotation	This policy specifies when a particular task is rotated among a group of employees.
Mandatory vacations	This policy specifies when an organization requires employees to take *X* number of consecutive days' vacation over the course of a year as part of their annual leave.
Onboarding	This policy takes place when a new employee is added to an organization and to its identity and access management system. It is associated with user training, federated identity management, and RBAC. *Offboarding* correlates to removing an employee from a federated identity management system.
Due diligence	This is the process of ensuring that IT infrastructure risks are known and managed.
Due care	This is the mitigation action that an organization takes to defend against the risks that have been uncovered during due diligence.
Due process	This is the principle that an organization must respect and safeguard the rights of personnel.
Clean desk policy	This policy directs employees as to how they should leave their working environment when they are not physically there to monitor it. For instance, they should leave the desk space free of any paperwork that may contain sensitive information. Also, they should log off any computing devices or phones.

Key Topic

Diversity of Training Techniques

The following sections describe diversity of training techniques. They include various methods and approaches to user security training such as gamification, capture the flag, phishing campaigns, phishing simulations, computer-based training (CBT), and role-based training.

User Education and Awareness Training

With so many possible organizational policies, employees need to be trained to at least get a basic understanding of them. Certain departments within an organization require more training than others. For example, human resources (HR) personnel need to understand many facets of the business and their corresponding policies, especially policies that affect personnel. HR people should be thoroughly trained in guidelines and enforcement. Sometimes the HR people train management and other employees on the various policies that those trainees are expected to enforce. In other cases, the trainer is an executive assistant or outside consultant.

Security awareness training is an ongoing process. Different organizations have varying types of security awareness training, and employees with different roles in the organization receive different types of training. This type of training is often coupled with the signing of a user agreement. The user, when signing this agreement, accepts and acknowledges specific rules of conduct, rules of behavior, and possibly the nondisclosure of any training (known as a nondisclosure agreement, or NDA).

NOTE Historically, the CompTIA exams themselves require that you sign an NDA. This means that you agree not to share any of the contents of the exams with anyone else.

All employees should be trained on ***personally identifiable information (PII)***. This is information used to uniquely identify, contact, or locate a person. This type of information could be a name, birthday, Social Security number, biometric information, and so on. Employees should know what identifies them to the organization and how to keep that information secret and safe from outsiders. Another key element of user education is the dissemination of the password policy. Employees should understand that passwords should be complex and that they should know the complexity requirements. They should also understand never to give out their password or ask for another person's password to any resource.

IT personnel should be trained on what to do in the case of account changes—for example, temporarily disabling the account of employees when they take a leave of absence or disabling the account (or deleting it, less common) of an employee who has been terminated. All IT personnel should be fluent in the organization's password policy, account lockout policy, and other user-related policies so that they can explain them to any other employees.

NOTE Typically, user security awareness training is done via computer-based training (CBT). This is normally done in the form of an interactive web page that provides the content for learning and interactive quizzes throughout to test the knowledge of the learner.

Some users might need to take additional privacy training, HIPAA training, or other types of security awareness training depending on the type of organization they work for. This user training might take the form of ***role-based training***, where the instructors and trainees act out the roles they might play, such as network administrator, security analyst, and so on. Instructors will often devise their training to take advantage of learning management systems and training metrics so that they can gauge the effectiveness of the training, validate compliance with policies, and analyze the security posture of the trainees in general.

User awareness training should also include interactive testing of the environment. For instance, many IT organizations employ ***phishing campaigns*** as part of their user awareness. This, of course, is not meant to be used in a malicious way. It is primarily used to test and educate personnel on how to handle phishing attempts. These attempts are usually carried out as ***phishing simulations*** by the InfoSec department of the organization. If employees are tricked by the phishing simulation, they are not penalized. They are simply directed to educational material specific to phishing techniques so that they can learn how to better handle the situation in the future.

Many InfoSec organizations also use ***gamification*** to educate their teams. This is typically in the form of some kind of ***capture the flag*** contest. A typical capture the flag training event within an organization would pit the organization's Red team against the Blue team. There are many opportunities for training of this type available online. Also, most major security conferences hold capture the flag contests. Such contests test the skills of security professionals in various ways. Some are very focused on Red team functions. Others are simply a *Jeopardy*-style question-and-answer challenge that allows participants to move forward based on answering the questions correctly. The idea is that this contest provides a way to challenge the participants and promotes a competitive environment for learning.

Key Topic Third-Party Risk Management

Before we begin this discussion of third-party risk management, we need to mention that the following information is *not* intended as legal advice. Before signing any contracts, an organization should strongly consider consulting with an attorney.

An organization often has in-depth policies concerning ***vendors***. Issues often occur because the level of agreement between the organization and the vendor was not clearly defined. A proper ***service-level agreement (SLA)*** that is analyzed by the organization carefully before signing can be helpful. A basic service contract is usually not enough; a service contract with an SLA will have a section within it that formally and clearly defines exactly what a vendor is responsible for and what the organization is responsible for—a demarcation point so to speak. It might also

define performance expectations and what the vendor will do if a failure of service occurs, timeframes for repair, backup plans, and so on. To benefit the organization, these points are usually legally binding and not informal. Consequently, it would benefit the organization to scrutinize the SLA before signing, and an organization's attorney should be involved in that process.

For instance, a company might use an ISP for its Internet connection. The customer will want to know what kind of fault-tolerant methods are on hand at the ISP and what kind of uptime to expect, which should be monitored by a network administrator. The SLA might have some sort of guarantee of measurable service that can be clearly defined—perhaps a minimum level of service and a target level of service. Before a company signs an SLA such as this, it is recommended that an attorney, the IT director, and other organizational management review the document carefully and make sure that it covers all the points required by the organization.

On a separate note, a ***business partnership agreement (BPA)*** is a type of contract that can establish the profits each partner will get, what responsibilities each partner will have, and exit strategies for partners. This often applies to ***supply chain*** and ***business partners***.

NOTE You will also see *BPA* stand for something else. An SLA that requires products and services over and over again is known as a blanket purchase agreement (BPA)—similar to a blanket order. These agreements are common in government contracts, but some organizations use them also. One issue to address is that there is some type of ending for the contract length. Some less-than-reputable cloud service providers will design open-ended BPAs. Try to avoid them.

Sometimes, multiple government agencies enter into a ***memorandum of understanding (MOU)***, or a letter of intent, in regard to a BPA; it could be that two agencies have a sort of convergence when it comes to ordering services.

Another type of agreement is the *interconnection security agreement (ISA)*. This agreement is established between two (or more) organizations that own and operate connected IT systems and data sets. Its purpose is to specifically document the technical and security requirements of the interconnection between organizations. This is the type of agreement you need in this scenario because the data is sensitive and the chief information officer (CIO) requires that there is a clear understanding of security controls to be implemented and agreed upon. As far as governing the security of data and systems, it is a more precise agreement than an SLA.

The ISA differs from the SLA, BPA, and MOU in the following ways:

- An SLA is a contract between a service provider and a customer that specifies the nature of the service to be provided and the level of service that the

provider will offer to the customer. It can be a very basic agreement, or it could also state the technical and performance parameters, but it will probably not include any specific security controls.

- A BPA does not have any inherent security planning in the way an ISA does.
- An MOU is not an agreement at all, but an understanding between two organizations or government agencies. It does not specify any security controls either. However, a memorandum of agreement (MOA) constitutes a legal agreement between two parties wishing to work together on a project but still does not detail any security controls.

Other important topics that should be addressed along with the service-level agreement (SLA) are ***end of life (EOL)*** and ***end of service life (EOSL)***. These topics define the product lifecycle and help your organization plan for future extensions and expansions of service levels. Also, if they are not adhered to, they could leave your environment vulnerable to compromise as well as lack of service when you have an outage. The EOL term typically defines the date when the product or service will no longer be sold or supported by the third party. EOSL is typically defined as the last day of service for the product, which means the third party no longer provides service if a failure occurs.

You don't need to get a postgraduate degree in business law, but it's a good idea to know these terms in case you need to communicate with the business people at your organization and to better understand the special contractual relationships between your organization and other organizations.

NOTE Measurement System Analysis (MSA) is a practical method of assessing the performance or usefulness of a specific measurement process.

Key Topic

Data Concepts

Sensitive data is information that can result in a loss of security or loss of advantage to a company if accessed by unauthorized persons.

Understanding Classification and Governance

Often, information is broken down into two groups: classified (which requires some level of security clearance) and nonclassified.

ISO/IEC 27002:2013 (which revises the older ISO/IEC 27002:2005) is a security standard that, among other things, can aid companies in classifying their data. Although you don't need to know the contents of that document for the Security+ exam, you should have a basic idea of how to classify information. For example, classification of data can be broken down as shown in Table 33-3.

Key Topic

Table 33-3 Data Sensitivity Classifications

Class	Description
Public information	Information available to anyone. Also referred to as unclassified or nonclassified.
Internal information	Information that is used internally by a company, but if it becomes public, no critical consequences result. This and the next three levels are known as *private* information. It might also be classified as *proprietary* information.
Confidential information	Information that can cause financial and operational loss to the company.
Secret information	Data that should never become public and is critical to the company.
Top secret information	The highest sensitivity of data; few people should have access, and security clearance may be necessary. Information is broken into sections on a need-to-know basis.

In this example, loss of public and internal information probably won't affect the company very much. However, unauthorized access, misuse, modification, or loss of confidential, secret, or top secret data can affect users' privacy, trade secrets, financials, and the general security of the company. By setting data roles such as data owner, data controller, data processor, data custodian/steward, and data privacy officer (DPO), and by classifying data and enforcing policies that govern who has access to what information, a company can limit its exposure to security threats. Different organizations classify data in various ways, but they are usually similar to Table 33-3. For example, you might also see the high, medium, and low classifications. Or, for instance, Red Hat Linux uses the Top Secret, Secret, and Confidential classifications (just as in Table 33-3), but considers everything else simply unclassified. All of these types of interpretations of data classifications are implementations of mandatory access control (MAC). It's the incorporation of these types of classifications that is a key element in the multilevel security of trusted operating systems (TOSs). Trusted operating systems such as Red Hat Linux, macOS, and HP-UX utilize multilevel security concepts such as these to meet government requirements.

NOTE *Data governance* is all of the data roles, processes, and controls that encompass the system inside of an organization that controls the authority over the actual data assets and how they are used. Its purpose is to ensure the quality and privacy of data throughout the data lifecycle.

Data Retention

A ***data retention*** policy states how long data must be stored by an organization. A retention policy should be in place for your log files, meaning they should be retained for future reference.

The log files can be secured in several ways:

- By employing the aforementioned backup methods.
- By setting permissions to the actual log file. File integrity is also important when securing log files. Encrypting the log files through the concept known as hashing is a good way to verify the integrity of these log files if they are moved and/or copied.
- By flat-out encrypting the entire contents of the file so that other users cannot view it. Encryption of any type does not sanitize the drive.

Storage retention and data retention usually manifest themselves as policies. For example, an organization might have a storage retention policy that states a hard drive must be kept in storage for a minimum of three years before being fully sanitized and/or destroyed. This policy is common in high-security environments where data is extremely confidential, or where auditing and other logging information must be kept for a specific amount of time.

Key Topic Credential Policies

The concepts discussed in this section are examples of what would normally be included in a ***personnel credential policy***. The most common type of authentication is the username/password combination. Usernames are usually based on a person's real name. Large organizations often use firstname.lastname as the standard naming convention (for example, ron.taylor@company.com) or first initial and last name (rtaylor@company.com). Smaller organizations might use the first name and last initial. The naming convention your organization chooses should be easy for you to implement without name confusion, and it should have the capability to be utilized for all systems on the network, including login, email, database, file access, and so on.

The password is either set by the user or created automatically for the user by an administrator or system. Figure 33-1 shows an example of an administrator-created password. However, in this case, the user is not blocked from changing the password (unless a policy was created for that purpose). Note that the second checkbox, User Cannot Change Password, is not selected. As an administrator, you also have the option to select User Must Change Password at Next Logon. A user would have to pick a password when first logging on to the domain, one that meets whatever

complexity requirements your network calls for. This, with the self-service password resetting (when users reset their own passwords at regular intervals), is necessary in larger networks to ease administration and increase security. The only caveat to this is account lockouts. Accounts that were locked by the system should only be unlocked by an administrator or system operator.

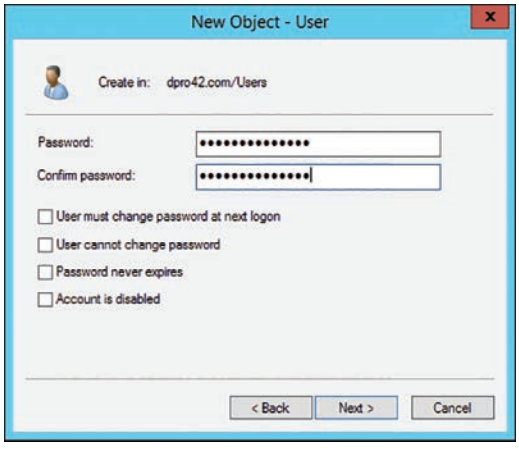


FIGURE 33-1 Password Phase of User Account Creation

At this point, it is common knowledge that a strong password is important for protecting a user account. Nowadays, many user accounts are compromised because of laziness—laziness on the part of the user for not configuring a strong password or lethargic complacency on the part of the administrator for not enforcing the use of strong passwords.

But what is a strong password? That depends on the organization you deal with, but generally it is broken down into a few easy-to-remember points. Passwords should comply with the following:

- Users no longer have to use special characters.
- Users should be able to use all characters.
- Copying and pasting passwords is acceptable.
- Password policies should not require employees to change passwords on a regular basis.

A weak password and vendor-set default password can lead to data breaches. Attackers use weak or known passwords to access accounts and then perform unauthorized copying, transfer, and retrieval of data from servers—collectively known as *data exfiltration*. Data exfiltration is also known as data extrusion, data exportation, or simply

stated, data theft. Whatever you call it, a complex password and properly configured password policies can help prevent it.

Changing your password at regular intervals is important as well. The general rule of thumb is to change your password as often as you change your toothbrush. However, because this is a subjective concept (to put it nicely!), many organizations have policies concerning passwords that we discuss in the next section. It might need to meet certain requirements, or be changed at regular intervals, and so forth.

The concepts in this example should also be applied to third-party credentials. For instance, if your employees have any cloud-based accounts, they should also follow the same password policies that you set forth for the internal systems whenever possible. This way, you can keep a consistent policy around credentials related to personnel.

Devices such as mobile phones or personal devices should also follow a standard policy created by the organization. For instance, many organizations require a passcode with a set number of characters to be used for access to the device. This information should be included in your organization's credential policy and enforced with a mobile device management system when possible. This way, you ensure that your employees' devices are secured from access by unauthorized users.

NOTE In many organizations, a mobile device management (MDM) system is utilized to implement mobile device use policies. This concept is discussed further in Chapter 21, "Implementing Secure Mobile Solutions."

When you're discussing policies for systems internal to your organization or devices that your organization owns, such as servers and laptops, it is also important to detail policies regarding how they should be deployed. For instance, ***service accounts*** should follow a specific credential policy. Service accounts often have higher-level permissions to enable a specific service. That means these accounts can cause more damage if they are compromised. They should always be configured using the ***least privilege*** approach. There should also be a service account for each service, instead of a common account used for many services. These are just a few of the recommendations that should be included in a credential policy in regard to service accounts.

Of course, ***administrator accounts*** or ***root accounts*** are the most dangerous accounts when it comes to compromise and must follow the strictest credential policy rules.



Organizational Policies

Next, let's dive into a discussion of organizational policies. Here you learn about change management, change control, and asset management concepts.

Change Management and Change Control

Change management is a structured way of changing the state of a computer system, network, or IT procedure. The idea behind this is that change is necessary, but that an organization should adapt to change and be knowledgeable of it. Any change that a person wants to make must be introduced to the heads of each department that it might affect. They must approve the change before it goes into effect. Before this happens, department managers will most likely make recommendations and/or give stipulations. When the necessary people have signed off on the change, it should be tested and then implemented. During implementation, it should be monitored and documented carefully. While change management is the overall process that the organization has put in place via their policies, ***change control*** is the specific part of that change management process that defines how a suggested change will be approved and in turn completed.

Because an IT infrastructure has so many interrelated parts and people, it is sometimes difficult for the left hand to know what the right hand is doing, or has done in the past. For example, after a network analysis, a network engineer might think that an unused interface on a firewall doesn't necessarily need to exist anymore. But does that engineer know this for sure? Who installed and configured the interface? When was it enabled? Was it ever used? Perhaps it is used only rarely by special customers making a connection to a screened subnet (DMZ); perhaps it is used with a honeynet; or maybe it is for future use or for testing purposes. It would be negligent for the network engineer to simply modify the firewall without at least asking around to find out whether the interface is necessary. More likely, there will be forms involved that require the network engineer to state the reason for change, and they must be signed by several other people before making the change. In general, this process will slow down progress, but in the long run it will help cover the network engineer. People were warned, and as long as the correct people involved have signed off on the procedure or technical change, the network engineer shouldn't have to worry. In a larger organization that complies with various certifications such as ISO 9001:2015, getting sign-off can be a complex task. IT people should have charts of personnel and department heads. There should also be current procedures in place that show who needs to be contacted in the case of a proposed change.

Asset Management

Asset management is a broad topic. For the purposes of organizational policy discussion, it is important to point out that understanding what is on your network and part of your computing environment is the first step in securing the environment. If you don't know that it is there, how can you defend it? Proper asset management should start with a strict policy for onboarding and offboarding devices that become

part of your environment. This means specifying how they are registered and activated and how they are later decommissioned.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 33-4 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 33-4 Key Topics for Chapter 33

Key Topic Element	Description	Page Number
Table 33-2	Policy Types	900
Section	Diversity of Training Techniques	900
Section	Third-Party Risk Management	902
Section	Data Concepts	904
Table 33-3	Data Sensitivity Classifications	905
Section	Credential Policies	906
Section	Organizational Policies	908

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

acceptable use policies (AUPs), separation of duties, job rotation, mandatory vacations, onboarding, user training, background checks, social media analysis, offboarding, personally identifiable information (PII), role-based training, phishing campaigns, phishing simulations, gamification, capture the flag, vendors, service-level agreement (SLA), business partnership agreement (BPA), supply chain, business partners, memorandum of understanding (MOU), end of life (EOL), end of service life (EOSL), data retention, personnel credential policy, service accounts, least privilege, administrator accounts, root accounts, change management, change control, asset management

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. What policy defines the rules that restrict how a computer, network, or other system may be used?
2. What is the security concept where more than one person is required to complete a particular task or operation?
3. Your company expects its employees to behave in a certain way. How could a description of this behavior be documented?
4. Employees are asked to sign a document that describes the methods of accessing a company's servers. What best describes this document?
5. One of the developers for your company asks you what to do before making a change to the code of a program's authentication. What process should you instruct this developer to follow?
6. As a network administrator, you are responsible for dealing with Internet service providers. You want to ensure that a provider guarantees end-to-end traffic performance. What is this known as?
7. What is considered information that available to anyone?
8. One of the accounting people is forced to change roles with another accounting person every three months. What is this an example of?
9. When it comes to security policies, what should HR personnel be trained in?



This chapter covers the following topics related to Objective 5.4 (Summarize risk management processes and concepts) of the CompTIA Security+ SY0-601 certification exam:

- Risk types
  - External
  - Internal
  - Legacy systems
  - Multiparty
  - IP theft
  - Software compliance/ licensing
- Risk management strategies
  - Acceptance
  - Avoidance
  - Transference
    - Cybersecurity insurance
  - Mitigation
- Risk analysis
  - Risk register
  - Risk matrix/heat map
  - Risk control assessment
  - Risk control self-assessment
  - Risk awareness
  - Inherent risk
  - Residual risk
  - Control risk
  - Risk appetite
  - Regulations that affect risk posture
  - Risk assessment types
    - Qualitative
    - Quantitative
  - Likelihood of occurrence
  - Impact
  - Asset value
  - Single loss expectancy (SLE)
  - Annualized loss expectancy (ALE)
  - Annualized rate of occurrence (ARO)
- Disasters
  - Environmental
  - Person-made
  - Internal vs. external
- Business impact analysis
  - Recovery time objective (RTO)
  - Recovery point objective (RPO)
  - Mean time to repair (MTTR)
  - Mean time between failures (MTBF)
  - Functional recovery plans
  - Single point of failure
  - Disaster recovery plan (DRP)
  - Mission essential functions
  - Identification of critical systems
  - Site risk assessment

Summarizing Risk Management Processes and Concepts

As it relates to computer security, a risk is the possibility of a malicious attack or other threat causing damage or downtime to a computer system. Generally, this is done by exploiting vulnerabilities in a computer system or network. The more vulnerability, the more risk. Organizations should be extremely interested in managing vulnerabilities and thereby managing risk. Risk management can be defined as the identification, assessment, and prioritization of risks, and the mitigating and monitoring of those risks. Specifically, when we talk about computer hardware and software, risk management is also known as information assurance (IA). In this chapter we start by discussing the different risk types and then move on to understanding risk management strategies and risk analysis. The chapter finishes with a discussion of disasters and business impact analysis.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 34-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 34-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Risk Types	1–2
Risk Management Strategies	3–5
Risk Analysis	6–8
Disaster Analysis	9
Business Impact Analysis	10–11

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What type of risk would a nation state threat actor be considered to a private corporation?
   a. External
   b. Internal
   c. Multilateral
   d. None of these answers are correct.

2. Your organization has just hired a new employee. However, the employee is really a spy looking to steal your intellectual property. Which type of risk would this be considered?
   a. External
   b. Internal
   c. Multivendor
   d. All of these answers are correct.

3. Which of the following is defined as the identification, assessment, and prioritization of risks, and the mitigating and monitoring of those risks?
   a. Risk management
   b. Risk mitigation
   c. Risk tolerance
   d. None of these answers are correct.

4. Which of the following usually entails not carrying out a proposed plan because the risk factor is too great?
   a. Risk transference
   b. Risk assessment
   c. Risk mitigation
   d. Risk avoidance

5. Which of the following is an example of risk transference?
   a. Risk appetite
   b. Risk avoidance
   c. Cybersecurity insurance
   d. All of these answers are correct.

6. Which of the following is the attempt to determine the number of threats or hazards that could possibly occur in a given amount of time to your computers and networks?
   a. Risk avoidance
   b. Risk assessment
   c. Risk mitigation
   d. Risk acceptance

7. _____ risk assessment is an assessment that assigns numeric values to the probability of a risk and the impact it can have on the system or network.
   a. Quantitative
   b. Qualitative
   c. Measured
   d. None of these answers are correct.

8. _____ is the total reduction or elimination of a risk.
   a. Risk assessment
   b. Risk transference
   c. Risk mitigation
   d. Risk acceptance

9. Which type of disaster can be defined as being caused by the influence of humans?
   a. Person-made
   b. Environmental
   c. Weather-made
   d. None of these answers are correct.

10. _____ defines the average number of failures per million hours of operation for a product in question.
    a. Mean time between failures (MTBF)
    b. Mean time to repair (MTTR)
    c. Mean time between recovery (MTBR)
    d. None of these answers are correct

11. _____ is the measured period of time between failures of a system.
    a. Mean time between failures (MTBF)
    b. Mean time to failure (MTTF)
    c. Mean time between recovery (MTBR)
    d. None of these answers are correct.

Foundation Topics

Key Topic Risk Types

According to NIST SP 800-137, the definition of risk is as follows:

> A measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence. [Note: Information system-related security risks are those risks that arise from the loss of confidentiality, integrity, or availability of information or information systems and reflect the potential adverse impacts to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, and the Nation. Adverse impacts to the Nation include, for example, compromises to information systems that support critical infrastructure applications or are paramount to government continuity of operations as defined by the Department of Homeland Security.]

Risks to your organization or environment can come in many shapes and forms. Let's first talk about the primary concern of most organizations: ***external risk***. This is, of course, the biggest concern to most because it is risk that comes from an external entity who could have many different motivations and therefore targets inside your organization. Many imagine an external "hacker" as someone sitting in a basement hammering away at the keyboard in front of 10 different monitors in a black hoodie. That, of course, is usually not the case. External risk most likely is from an organized threat actor or an organization of threat actors who have various motivations. To carry out their attacks, they will use many different methods, many of which we have discussed in depth in other sections of this book. The primary goal of external attackers is to gain access to your organization's computing environment, gain a foothold, and keep it as long as possible to carry out their objectives, whatever they may be.

Of course, what many organizations tend to overlook are the ***internal risks***. The majority of internal risks stems from employees or those internal to the organization such as contractors. While the motivation of external cybercriminals may be to gain access and keep access, the internal threat actor already has access to the organization's environment. This person's motivations are usually different from those of external threat actors—although the goal may be the same in the end. Most internal attacks result in the exfiltration or destruction of sensitive data that belongs to the organization. The ***theft of intellectual property*** is a primary goal of internal and external threat actors.

Let's talk about the risks that exist within the organization that could lead to internal or external threat actors obtaining their goal. As we have discussed, external threat actors want to maintain access that they have obtained. To do this, they often use the technique of pivoting and scanning the internal environment that they now have access to. An easy target for them to gain another foothold would be a legacy system on the network that is out of date and contains vulnerabilities that allow an additional attack to be carried out. Many organizations take the approach of a hard outer shell to secure their network and allow for a softer internal attack surface. Often this is caused by a lack of ***software compliance/licensing*** processes within the organization. The belief in these organizations is that the internal ***legacy systems*** are not at risk because they are behind a firewall and not directly accessible from the Internet. As you know, this is a misconception. As soon as your external network is breached, all of those internal legacy systems become prime targets. Of course, let's not forget that in the cybersecurity world things are not exactly black and white. Some risks would be considered ***multiparty***. External threat actors could be working with internal threat actors. Also, external threat actors could have access to an asset such as a contractor laptop that then gives them that internal access they are looking to obtain.



Risk Management Strategies

In this section we dive deeper into the details of risk management strategies. Organizations usually employ one of the four following general strategies when managing a particular risk:

- Transfer the risk to another organization or third party.
- Avoid the risk.
- Mitigate the risk.
- Accept some or all of the consequences of a risk.

It is possible to transfer *some* risk to a third party. An example of ***risk transference*** (also known as risk sharing) would be an organization that purchases ***cybersecurity insurance*** for a group of servers in a data center. The organization still takes on the risk of losing data in the case of server failure, theft, and disaster but transfers the risk of losing the money those servers are worth in case they are lost.

Some organizations opt to avoid risk. ***Risk avoidance*** usually entails not carrying out a proposed plan because the risk factor is too great. An example of risk avoidance: If a high-profile organization decided not to implement a new and controversial website based on its belief that too many attackers would attempt to exploit it.

However, the most common goal of risk management is to *mitigate* all risk to a level acceptable to the organization. It is impossible to eliminate all risk, but it should be mitigated as much as possible within reason. Usually, budgeting and IT resources dictate the level of ***risk mitigation*** and what kinds of deterrents can be put in place. For example, installing antivirus/firewall software on every client computer is common; most companies do this. However, installing a high-end, hardware-based firewall at every computer is not common; although this method would probably make for a secure network, the amount of money and administration needed to implement that solution would make it unacceptable.

Most organizations are willing to accept a certain amount of risk. This is ***risk acceptance***, also known as risk retention. Sometimes, vulnerabilities that would otherwise be mitigated by the implementation of expensive solutions are instead dealt with when and if they are exploited. IT budgeting and resource management are big factors when it comes to these risk management decisions.

After the risk transference, risk avoidance, and risk mitigation techniques have been implemented, an organization is left with a certain amount of ***residual risk***—the risk left over after a detailed security plan and disaster recovery plan have been implemented. There is always risk because a company cannot possibly foresee every future event, nor can it secure against every single threat. Senior management as a collective whole is ultimately responsible for deciding how much residual risk there will be in a company's infrastructure and how much risk there will be to the company's data. Often, no one person will be in charge of this level, but it will be decided on as a group.

There are many different types of risks to computers and computer networks. Of course, before you can decide what to do about particular risks, you need to assess what those risks are.

Key Topic Risk Analysis

Risk assessment is the attempt to determine the number of threats or hazards that could possibly occur in a given amount of time to your computers and networks. When you assess risks, they are often recognized threats, but risk assessment can also take into account new types of threats that might occur. When risk has been assessed, it can be mitigated up until the point at which the organization will accept any residual risk. Generally, risk assessments follow a particular order, for example:

Step 1. Identify the organization's assets.

Step 2. Identify vulnerabilities.

Step 3. Identify threats and ***threat*** likelihood.

Step 4. Identify potential monetary impact.

The fourth step is also known as ***impact assessment***. At this point, you determine the potential monetary costs related to a threat.

An excellent tool to create during your risk assessment is a ***risk register***, also known as a risk log, which helps track issues and address problems as they occur. After the initial risk assessment, you, as security administrator, will continue to use and refer to the risk register. It can be a great tool for just about any organization but can be of more value to certain types of organizations, such as manufacturers that utilize a supply chain. In this case, the organization would want to implement a specialized type of risk management called *supply chain risk management (SCRM)*. In this type, the organization collaborates with suppliers and distributors to analyze and reduce risk. One approach that an organization can take to identify risks and controls is the application of a ***risk control assessment***. This assessment can be completed by a third party or by an internal team. The internal team can complete a ***risk control self-assessment*** (RCSA). This process is often used by financial institutions that need to meet regulatory compliances. A great approach to visually representing the results of a risk assessment is the use of a ***risk matrix/heat map***. See Figure 34-1 for an example.

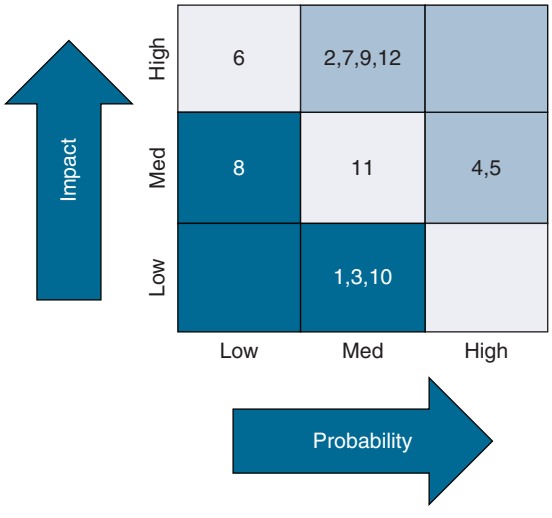


FIGURE 34-1 Risk Matrix/Heat Map

Table 34-2 summarizes common terms associated with risk analysis.

Table 34-2 Common Risk Terms

Term	Description
Risk appetite	The types and amount of risk, on a broad level, an organization is willing to accept in its pursuit of value.
Inherent risk	The representation of the level of risk an organization would experience if the correct mitigation was not in place.
Residual risk	The risk left over after a detailed security plan and disaster recovery plan have been implemented.
Control risk	The risk that a control that is in place may not detect or may fail to protect the environment.
Risk awareness	The ability of an organization to identify risks before they become a threat. The overall preparedness of an organization to mitigate risk.
Risk mitigation	NIST defines risk mitigation as "Prioritizing, evaluating, and implementing the appropriate risk-reducing controls/countermeasures recommended from the risk management process."

The two most common risk assessment methods are qualitative and quantitative.

Qualitative Risk Assessment

Qualitative risk assessment is an assessment that assigns numeric values to the probability of a risk and the impact it can have on the system or network. Unlike its counterpart, quantitative risk assessment, it does not assign monetary values to assets or possible losses. It is the easier, quicker, and cheaper way to assess risk but cannot assign ***asset value*** or give a total for possible monetary loss.

With this method, ranges can be assigned, for example, 1 to 10 or 1 to 100. The higher the number, the higher the probability of risk, or the greater the ***impact*** on the system. As a basic example, a computer without antivirus software that is connected to the Internet will most likely have a high probability of risk; it will also most likely have a great impact on the system. You could assign the number 99 as the probability of risk. You are not sure exactly when it will happen but are 99 percent sure that it will happen at some point. Next, you could assign the number 90 out of 100 as the impact of the risk. This number implies a heavy impact; probably either the system has crashed or has been rendered unusable at some point. There is a 10 percent chance that the system will remain usable, but it is unlikely. Finally, you multiply the two numbers together to find out the qualitative risk: $99 \times 90 = 8910$. That's 8910 out of a possible 10,000, which is a high level of risk. The way to mitigate risk in this example would be to install antivirus software and verify that it is configured to auto-update. By assigning these types of qualitative values to various risks, you can make comparisons from one risk to another and get a better idea of what needs to be mitigated and what doesn't.

The main issue with this type of risk assessment is that it is difficult to place an exact value on many types of risks. The type of qualitative system varies from organization to organization, even from person to person; it is a common source of debate as well. This makes qualitative risk assessments more descriptive than truly measurable. However, by relying on group surveys, company history, and personal experience, you can get a basic idea of the risk involved.

Quantitative Risk Assessment

Quantitative risk assessment measures risk by using exact monetary values. It attempts to give an expected yearly loss in dollars for any given risk. It also defines ***asset values*** to servers, routers, and other network equipment.

Three values are used when making quantitative risk calculations:

- ***Single loss expectancy (SLE)*:** The loss of value in dollars based on a single incident.
- ***Annualized rate of occurrence (ARO)*:** The number of times per year that the specific incident occurs.
- ***Annualized loss expectancy (ALE)*:** The total loss in dollars per year due to a specific incident. The incident might happen once or more than once; either way, this number is the total loss in dollars for that particular type of incident. It is computed with the following calculation:

  $SLE \times ARO = ALE$

So, for example, suppose you wanted to find out how much an e-commerce web server's downtime would cost the company per year. You would need some additional information such as the average web server downtime in minutes and the number of times this occurs per year. You also would need to know the average sale amount in dollars and how many sales are made per minute on this e-commerce web server. This information can be deduced by using accounting reports and by further security analysis of the web server, which we discuss later. For now, let's just say that over the past year the web server failed seven times. The average downtime for each failure was 45 minutes. That equals a total of 315 minutes of downtime per year, close to 99.9 percent uptime. (The more years you can measure, the better the estimate will be.) Now let's say that this web server processes an average of 10 orders per minute with average revenue of $35. That means that $350 of revenue comes in per minute. As mentioned, a single downtime averages 45 minutes, corresponding to a $15,750 loss per occurrence. So, the SLE is $15,750. Ouch! Some salespeople are going to be unhappy with your 99.9 percent uptime! But you're not done. You can calculate the annualized loss expectancy (ALE) by multiplying the SLE ($15,750) by the annualized rate of occurrence (ARO). Because the web server failed seven times

last year, the SLE × ARO would be $15,750 × 7, which equals $110,250 (the ALE). This example is shown in Table 34-3.

Key Topic

Table 34-3 Quantitative Risk Assessment Example

SLE	ARO	ALE
$15,750	7	$110,250
Revenue lost due to each web server failure	Total web server failures over the past year	Total loss due to web server failure per year

Apparently, you need to increase the uptime of the e-commerce web server! Many organizations demand 99.99 percent or even 99.999 percent uptime; 99.999 percent uptime means that the server will have only 5 minutes of downtime over the entire course of the year. Of course, to accomplish this, you first need to scrutinize the server to see precisely why it fails so often. What exactly are the vulnerabilities of the web server? Which ones were exploited? Which threats exploited those vulnerabilities? By exploring the server's logs, configurations, and policies, and by using security tools, you can discern exactly why this happens so often. However, this analysis should be done carefully because the server does so much business for the company.

It isn't possible to assign a specific ALE to incidents that will happen in the future, so new technologies should be monitored carefully. Any failures should be documented thoroughly. For example, a spreadsheet could be maintained that contains the various technologies your organization uses; their failure history; their SLE, ARO, and ALE; and mitigation techniques that you have employed, and when they were implemented.

Table 34-4 compares the different aspects of quantitative and qualitative risk.

Key Topic

Table 34-4 Risk Assessment Types

Risk Assessment Type	Description	Key Points
Qualitative risk assessment	Assigns numeric values to the probability of a risk, and the impact it can have on the system or network.	Numbers are arbitrary. Examples: 1–10 or 1–100.
Quantitative risk assessment	Measures risk by using exact monetary values. It attempts to give an expected yearly loss in dollars for any given risk.	Values are specific monetary amounts. SLE × ARO = ALE MTBF can be used for additional data.

NOTE Most organizations within the medical, pharmaceutical, and banking industries use quantitative risk assessments; they need to have specific monetary numbers to measure risk. Taking this one step further, many banking institutions adhere to the recommendations within the Basel I, II, and III Accords. These recommended standards describe how much capital a bank should put aside to aid with financial and operational risks if they occur.

Key Topic

Disaster Analysis

No matter how much redundancy you implement, there is always a chance that a disaster could arise. A disaster could be the loss of data on a server, a fire in a server room, or the catastrophic loss of access to an organization's building. To prepare for these events, you should design a disaster recovery plan, but with the thought in mind that redundancy and fault tolerance can defend against most "disasters." The best administrator is the one who avoids disaster and, in the rare case that it does happen, has a plan in place to recover quickly from it.

Before you can plan for disasters, you need to define exactly what disasters are possible and list them in order starting with the most probable. This step sounds a bit morbid, but it's necessary to ensure the long-term welfare of your organization.

Disasters can be divided into two categories: ***environmental*** and ***person-made***. They can also be looked at from the perspective of internal versus external. Some of the disasters that could render your server room inoperable include the following:

- **Fire:** Fire is probably the number one planned-for disaster. This is partially because most municipalities require some sort of fire suppression system, as well as the fact that most organizations' policies define the usage of a proper fire suppression system. The three main types of fire extinguishers include A (for ash fires), B (for gas and other flammable liquid fires), and C (for electrical fires). Unfortunately, these and the standard sprinkler system in the rest of the building are not adequate for a server room. If there were a fire, the material from the fire extinguisher or the water from the sprinkler system would damage the equipment, making the disaster even worse! Instead, a server room should be equipped with a proper system of its own such as DuPont FM-200. This system uses a large tank that stores a clean agent fire extinguishant that is sprayed from one or more nozzles in the ceiling of the server room. It can put out fires of all types in seconds. A product such as this can be used safely when people are present; however, most systems also employ a *very* loud alarm that tells all personnel to leave the server room. It is wise to run through

several fire suppression alarm tests and fire drills, ensuring that the alarm will sound when necessary and that personnel know what do to when the alarm sounds. For example, escape plans should be posted, and battery-backup exit signs should be installed in various locations throughout the building so that employees know the quickest escape route in the case of a fire. Fire drills (and other safety drills) should be performed periodically so that the organization can analyze the security posture of their safety plan.

- **Flood:** The best way to avoid server room damage in the case of a flood is to locate the server room on the first floor or higher, not in a basement. There's not much you can do about the location of a building, but if it is in a flood zone, it makes the use of a warm or hot site that much more imperative. And a server room could also be flooded by other things such as boilers. The room should not be adjacent to, or on the same floor as, a boiler room. It should also be located away from other water sources such as bathrooms and any sprinkler systems. The server room should be thought of three-dimensionally; the floors, walls, and ceiling should be analyzed and protected. Some server rooms are designed to be a room within a room and might have drainage installed as well.
- **Long-term power loss:** Short-term power loss should be countered by the UPS, but long-term power loss requires a backup generator and possibly a redundant site.
- **Theft and malicious attack:** Theft and malicious attack can also cause a disaster, if the right data is stolen. Physical security such as door locks/access systems and video cameras should be implemented to avoid this. Servers should be cable-locked to their server racks, and removable hard drives (if any are used) should have key access. Not only do you, as security administrator, have the task of writing policies and procedures that govern the security of server rooms and data centers, but you will often have the task of *enforcing* those policies—meaning muscle in the form of security guards, and dual-class technician/guards—or by otherwise having the right to terminate employees as needed, contact and work with the authorities, and so on.
- **Loss of building:** Temporary loss of the building due to gas leak, malicious attack, inaccessibility due to crime scene investigation, or natural event will require personnel to access a redundant site. Your server room should have as much data archived as possible, and the redundant site should be warm enough to keep business running. A plan should be in place as to how data will be restored at the redundant site and how the network will be made functional.

Key Topic Business Impact Analysis

Next, we discuss business impact analysis. We start by covering recovery time objectives and recovery point objectives, continuing on with a description of mean time to repair and mean time between failures. From there, we cover functional recovery plans, single point of failure, and a disaster recovery plan. We conclude with a look it mission-essential functions, identification of critical systems, and site risk assessment.

Although it's impossible to predict the future accurately, it can be quantified on an average basis using concepts such as ***mean time between failures (MTBF)***. This term deals with reliability. It defines the average number of failures per million hours of operation for a product in question. This number is based on historical baselines among various customers who use the product. It can be very helpful when making quantitative assessments.

NOTE Another way of describing MTBF is called failure in time (FIT), which is the number of failures per *billion* hours of operation.

You should know two other terms related to MTBF: ***mean time to repair (MTTR)***, which is the time needed to repair a failed device, and ***mean time to failure (MTTF)***, which is a basic measure of reliability for devices that cannot be repaired. All three of these concepts should also be considered when creating a ***disaster recovery plan (DRP)***.

When an environment is planned properly, it can withstand most failures barring total disaster using the following redundancy precautions:

- Redundant power in the form of power supplies, UPSs, and backup generators
- Redundant data, servers, ISPs, and sites

The whole concept revolves around single points of failure. A ***single point of failure*** is an element, object, or part of a system that, if it fails, causes the whole system to fail. By implementing redundancy, you can bypass just about any single point of failure.

There are two methods to combating single points of failure. The first is to use redundancy. If employed properly, redundancy keeps a system running with no downtime. However, this solution can be pricey, and we all know there is only so much IT budget to go around. So, the alternative is to make sure you have plenty of spare parts lying around. This is a good method if your network and systems are not time-critical. Installing spare parts often requires you to shut down the server or a portion of a network. If this risk is not acceptable to an organization, you'll have to

find the cheapest redundant solutions available. Research is key, but don't be fooled by the hype: sometimes the simplest sounding solutions are the best.

Here's the scenario. Your server room has the following powered equipment:

- Nine servers
- Two Microsoft domain controllers (DCs)
- One DNS server
- Two file servers
- One database server
- Two web servers (which second as FTP servers)
- One mail server
- Five 48-port switches
- One master switch
- Three routers
- Two CSU/DSUs
- One PBX
- Two client workstations (for remote server access without having to work directly at the server) within the server room

It appears that there is already some redundancy in place in this server room. For example, there are two domain controllers. One of them has a copy of Active Directory and acts as a secondary DC in case the first one fails. There are also two web servers, one ready to take over for the other if the primary one fails. This type of redundancy is known as failover redundancy. The secondary system is inactive until the first one fails. Also, two client workstations are used to remotely control the servers; if one fails, another one is available.

Otherwise, the rest of the servers and other pieces of equipment are one-offs—single instances in need of something to prevent failure. There are a lot of them, so you truly need to *redundacize*. Hey, it's a word if IT people use it! It's the detailed approach to preparing for problems that can arise in a system that will make for a good IT contingency plan. Try to envision the various upcoming redundancy methods used with each of the items listed previously in the fictitious server room.

But before we get into some hard-core redundancy, let's discuss the terms *fail-open* and *fail-closed*. *Fail-open* means that if a portion of a system fails, the rest of the system will still be available or "open." *Fail-closed* means that if a portion of a system

fails, the entire system will become inaccessible or simply shut down. Depending on the level of security your organization requires, you might have a mixture of fail-open and fail-closed systems. In the previous server room example, there is a DNS server and a database server. Let's say that the DNS server forwards information to several different zones and that one of those zones fails for one reason or another. You might decide that it is more beneficial to the network to have the rest of the DNS server continue to operate and service the rest of the zones instead of shutting down completely, so you would want the DNS server to fail-open. However, the database server might have confidential information that you cannot afford to lose, so if one service or component of the database server fails, you might opt to have the database server stop servicing requests altogether, or in other words, to fail-closed. Another example would be a firewall/router. If the firewall portion of the device failed, you would probably want the device to fail-closed. Even though the network connectivity could still function, you probably wouldn't want it to because there is no firewall protection. Your solution depends on the level of security you require and the risk that can be associated with devices that fail-open. It also depends on whether the server or device has a redundancy associated with it. If the DNS server mentioned previously has a secondary redundant DNS server that is always up and running and ready to take requests at a moment's notice, you might opt to instead configure the first DNS server to fail-closed and let the secondary DNS server take over entirely.

Disaster Recovery Planning

Disaster recovery plans should include information regarding redundancy, such as sites and backup, but should not include information that deals with the day-to-day operations of an organization, such as updating computers, patch management, monitoring and audits, and so on. It is important to include only what is necessary in a disaster recovery plan. Too much information can make it difficult to use when a disaster does strike.

Although not an exhaustive set, the following written disaster recovery policies, procedures, and information should be part of your disaster recovery plan:

- **Contact information:** You should identify the people or resources to contact if a disaster occurs and how employees will contact the organization.
- **Impact determination:** This procedure determines a disaster's full impact on the organization. It includes an evaluation of assets lost and the cost to replace those assets.
- **Functional recovery plan:** This plan is based on the determination of disaster impact. It will have many permutations depending on the type of disaster.

The recovery plan includes an estimated time to complete recovery and a set of steps defining the order of what will be recovered and when. It might also include an after action report (AAR), which is a formal document designed to determine the effectiveness of a recovery plan in the case that it was implemented.

- **Business continuity plan:** This plan defines how the business will continue to operate if a disaster occurs; the BCP is often carried out by a team of individuals. A BCP is also referred to as a continuity of operations plan (COOP). Over the years, BCPs have become much more important, and depending on the organization, a BCP might actually encompass the entire DRP. It also comprises business impact analysis—the examination of critical versus non-critical functions. These functions are assigned two different values or metrics: ***recovery time objective (RTO)***, the acceptable amount of time to restore a function (for example, the time required for a service to be restored after a disaster), and ***recovery point objective (RPO)***, the acceptable latency of data, or the maximum tolerable time that data can remain inaccessible after a disaster. It's impossible to foresee exactly how long it will take to restore service after a disaster, but with the use of proper archival, hot/warm/cold sites, and redundant systems, a general timeframe can be laid out, and an organization will be able to decide on a maximum timeframe to get data back online. This, in effect, is IT contingency planning (ITCP).

Some organizations have a continuity of operation planning group or crisis management group that meets every so often to discuss the BCP. Instead of running full-scale drills, they might run through tabletop exercises, where a talk-through of simulated disasters (in real time) is performed—a sort of role playing, if you will. This approach can save time and be less disruptive to employees, but it is more than just a read-through of the BCP. It can help identify critical systems and mission-essential functions of the organization's network as well as failover functionality and alternate processing sites. It can also aid in assessing the impact of a potential disaster on privacy, property, finance, the reputation of the company, and most importantly, life itself.

- **Copies of agreements:** Copies of any agreements with vendors of redundant sites, ISPs, building management, and so on should be stored with the DR plan.
- **Disaster recovery drills and exercises:** Employees should be drilled on what to do if a disaster occurs. These exercises should be written out step by step and should conform to safety standards.
- **Hierarchical list of critical systems and critical data:** This list includes all the ***mission-essential functions*** and ***identification of critical systems*** necessary

for business operations: domain controllers, firewalls, switches, DNS servers, file servers, web servers, and so on. They should be listed by priority. Systems such as client computers, test computers, and training systems would be last on the list or not listed at all. You should also include (somewhere in the DRP) some geographic considerations. For example, are there offsite backups or virtualization in place? What is the physical distance to those backups and virtual machines? And, are there legal implications? For instance, are there data sovereignty implications—meaning, will it be difficult to gain access to data and VMs stored in a different country based on the laws of that country? For each disaster recovery site, a ***site risk assessment*** should be completed to determine the risk based on the actual site-specific conditions. For instance, the geographical location of the DR site might raise additional risks that other sites do not face.

Generally, the chief security officer (CSO) or other high-level executive will be in charge of DR planning, often with the help of the information systems security officer (ISSO); however, who is in charge depends on the size of the organization and the types of management involved. That said, any size organization can benefit from proper DR planning. This information should be accessible at the company site, and a copy should be stored offsite as well. If your organization conforms to special compliance rules, you should consult them when designing a DR plan. Depending on the type of organization, yet other items might go into your DR plan.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 34-5 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 34-5 Key Topics for Chapter 34

Key Topic Element	Description	Page Number
Section	Risk Types	917
Section	Risk Management Strategies	918
Section	Risk Analysis	919

Key Topic Element	Description	Page Number
Table 34-3	Quantitative Risk Assessment Example	923
Table 34-4	Risk Assessment Types	923
Section	Disaster Analysis	924
Section	Business Impact Analysis	926

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

external risk, internal risk, theft of intellectual property, software compliance/licensing, legacy systems, multiparty, risk management, information assurance (IA), risk transference, cybersecurity insurance, risk avoidance, risk mitigation, risk acceptance, residual risk, risk assessment, threat likelihood, impact assessment, risk register, risk control assessment, risk control self-assessment, risk matrix/heat map, risk appetite, inherent risk, residual risk, risk awareness, risk mitigation, qualitative risk assessment, asset value, impact, quantitative risk assessment, asset values, single loss expectancy (SLE), annualized rate of occurrence (ARO), annualized loss expectancy (ALE), environmental disaster, person-made disaster, mean time between failures (MTBF), mean time to repair (MTTR), mean time to failure (MTTF), disaster recovery plan (DRP), single point of failure, recovery time objective (RTO), recovery point objective (RPO), mission-essential functions, identification of critical systems, site risk assessment

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. Which type of plan is based on the determination of disaster impact?
2. _____is the time required for a service to be restored after a disaster.
3. What procedure is used to determine a disaster's full impact on the organization?
4. What is considered the risk left over after a detailed security plan and disaster recovery plan have been implemented?
5. What is considered an element, object, or part of a system that, if it fails, causes the whole system to fail?

6. _____ defines the average number of failures per million hours of operation for a product in question.
7. Which type of assessment measures risk by using exact monetary values?
8. What term is used when risk is reduced or eliminated altogether?
9. Which type of assessment assigns numeric values to the probability of a risk and the impact it can have on the system or network?
10. What is the attempt to determine the number of threats or hazards that could possibly occur in a given amount of time to your computers and networks?

This page intentionally left blank



This chapter covers the following topics related to Objective 5.5 (Explain privacy and sensitive data concepts in relation to security) of the CompTIA Security+ SY0-601 certification exam:

- Organizational consequences of privacy and data breaches
  - Reputation damage
  - Identity theft
  - Fines
  - IP theft
- Notifications of breaches
  - Escalation
  - Public notifications and disclosures
- Data types (Classifications)
  - Public
  - Private
  - Sensitive
  - Confidential
  - Critical
  - Proprietary
  - Personally identifiable information (PII)
  - Health information
  - Financial information
  - Government data
  - Customer data
- Privacy enhancing technologies
  - Data minimization
  - Data masking
  - Tokenization
  - Anonymization
  - Pseudo-anonymization
- Roles and responsibilities
  - Data owner
  - Data controller
  - Data processor
  - Data custodian/steward
  - Data protection officer (DPO)
- Information life cycle
- Impact assessment
- Terms of agreement
- Privacy notice

CHAPTER 35

Understanding Privacy and Sensitive Data Concepts in Relation to Security

This chapter covers organizational consequences of privacy and data breaches, including reputation damage, identity theft, fines, and IP theft. We also discuss notifications of breaches, including escalation and public notification and disclosures. The chapter continues with an overview of data types and classifications such as public, private, sensitive, confidential, critical, and proprietary. You also learn about personally identifiable information (PII), health information, financial information, government data, and customer data. In addition, you learn about privacy enhancing technologies, including data minimization, data masking, tokenization, anonymization, and pseudo-anonymization. This chapter also covers the roles and responsibilities such as data owner, data controller, data processor, data custodian/steward, and data protection officer (DPO). The chapter finishes with an explanation of the information lifecycle, impact assessment, terms of agreement, and privacy notice.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the "Chapter Review Activities" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 35-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Review Questions."

Table 35-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Organizational Consequences of Privacy and Data Breaches	1
Notifications of Breaches	2
Data Types and Asset Classification	3–6
Privacy Enhancing Technologies	7, 8
Roles and Responsibilities	9
Information Lifecycle	10
Impact Assessment	11
Terms of Agreement	12
Privacy Notice	13

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What term is used to describe the crown jewels of an organization?
   a. Intellectual property
   b. Important data
   c. Sensitive data
   d. None of these answers are correct.

2. Which organization requires that any publicly traded company provide a public notification and disclosure of a data breach?
   a. FBI
   b. SEC
   c. CIA
   d. All of these answers are correct.

3. Which type of data classification would cause grave damage to national security?
   a. Top secret
   b. Confidential
   c. Unclassified
   d. None of these answers are correct.

4. Which type of data classification would cause no damage to national security?
   a. Secret
   b. Confidential
   c. Top secret
   d. Unclassified

5. Unauthorized access to this type of data could cause severe damage to an organization?
   a. Nonsensitive
   b. Sensitive
   c. Public
   d. Private

6. What information can be used to distinguish or trace an individual's identity?
   a. PHI
   b. PII
   c. HIPAA
   d. None of these answers are correct.

7. _____ is a term used to explain the concept of reducing the amount of personal information consumed by online entities.
   a. Data masking
   b. Data minimization
   c. Data usage
   d. Data processing

8. Which of the following is used to obfuscate sensitive data?
   a. Data minimization
   b. Data processing
   c. Data masking
   d. Data control

9. Which role is usually part of the management team and maintains ownership of a subset of data?
   a. Data custodian
   b. Data owner
   c. Data processor
   d. All of these answers are correct.

10. During which phase of the information lifecycle is data obtained?
    a. Disposal
    b. Storage
    c. Usage
    d. Collection

11. Which process that is required by GDPR involves identifying the risk of data compromise to an individual?
    a. Impact assessment
    b. Data control
    c. Terms of agreement
    d. All of these answers are correct.

12. Which concept adopted and required by GDPR helps protect the personal information of an individual?
    a. Data classification
    b. Data inspection
    c. Data collection
    d. Terms of agreement

13. Which document must be provided to individuals defining how their data will be used?
    a. GDPR
    b. Data collection
    c. Terms of agreement
    d. Privacy notice

Foundation Topics

Key Topic Organizational Consequences of Privacy and Data Breaches

When it comes to the consequences of privacy and data breaches on an organization, a few factors are typically used to determine the actual impact on the organization. Those factors include reputation damage, ***identity theft***, ***fines***, and ***intellectual property (IP) theft***. Let's take a deeper look into what each factor really means.

Reputation damage can be attributed to any kind of breach; however, many times the types of breaches that result in reputation damage involve the loss of customer data. These types of compromises typically get more press and in turn become a catalyst for reputation damage. Of course, the impact also depends on the type of organization that incurred the breach and loss of data. For instance, if an organization that is responsible for safely storing personal information about its customers has a breach disclosing that information, the situation is more significant. The resulting outcome of this kind of data breach could be ***identity theft*** for its customers. This theft, of course, would significantly affect the company's brand reputation, which could take a long time to recover, if it ever does.

NOTE In ***identity theft***, your personal information is stolen and subsequently used to commit an act of fraud in your name. Many times this act affects credit cards, taxes, and medical records. This theft can, of course, cause severe damage to your credit and take a lot of time to restore.

Depending on the industry where the compromise takes place, it can also result in ***fines***. For instance, the health-care industry regulations fall under the Health Insurance Portability and Accountability Act of 1996 (HIPAA). Noncompliance with these regulations carries large ***fines*** and can even result in jail time for the individuals responsible.

Regardless of what industry you are responsible for protecting, the ultimate goal is to protect your intellectual property. This is, of course, the organization's crown jewels and can be in many different forms. For a company that develops software, the crown jewels are the source code of the product it is selling. Even a company that produces food products has intellectual property it is trying to protect, such as a secret recipe.

Notifications of Breaches

When a breach happens, specific procedures should be followed based on the regulations of the state where the organization does business. Security breach notification laws require the compromised organization to notify all individuals impacted by the data breach of the personal information that was disclosed. Each state has different legislation enacted to enforce this notification. The intention of such a notification is to provide the impacted individuals with enough information that they will be able to mitigate the possible risk imposed by the data breach. Some of the laws include a requirement for ***public notifications and disclosures***. For instance, the US Securities and Exchange Commission (SEC) requires that any publicly traded company provide a public notification and disclosure of a data breach. Many factors define the actual process of ***escalation*** when it comes to data breaches. As discussed, these can include the type of industry, the location of the company doing business, and other factors. An organization should take all these factors into consideration when developing an incident response plan. One of the first factors that should be taken into account is the actual data type.

Key Topic Data Types and Asset Classification

To protect an asset, an organization first needs to understand how important that asset is. For example, the unauthorized disclosure of the source code of a product might be more impactful on an organization than the disclosure of a public configuration guide. The first step in implementing an access control process is to classify assets or data based on the potential damage a breach to the confidentiality, integrity, or availability of that asset or data could cause.

This process is called asset or data classification, and there are several ways to classify assets. For example, military and governmental organizations commonly use the following classification definitions:

- **Top Secret:** Unauthorized access to top secret information would cause grave damage to national security.
- **Secret:** Unauthorized access to secret information would cause severe damage to national security.
- **Confidential:** Unauthorized access to confidential information would cause damage to national security.
- **Unclassified:** Unauthorized access to unclassified information would cause no damage to national security.

The commercial sector has more variety in the way data classification is done—more specifically, to the label used in the classification. Here are some commonly used classification labels in the commercial sector:

- ***Confidential* or *Proprietary*:** Unauthorized access to confidential or proprietary information could cause grave damage to the organization. Examples of information or assets that could receive this type of classification include source code and trade secrets.
- **Private:** Unauthorized access to private information could cause severe damage to the organization. Examples of information or assets that could receive this type of classification are human resource information (for example, employee salaries), medical records, and so on.
- ***Sensitive*:** Unauthorized access to sensitive information could cause some damage to the organization. Examples of information or assets that could receive this type of classification are internal team email, financial information, and so on.
- ***Public*:** Unauthorized access to public information does not cause any significant damage.
- ***Critical*:** This is data that is critical to the continued function of a business. Loss of this type of data would result in significant monetary loss.

Although the classification schema will differ from one company to another, it is important that all departments within a company use the schema consistently. For each label, there should be a clear definition identifying when that label should be applied and what damage would be caused by unauthorized access. Because the classification of data may also be related to specific times or other contextual factors, the asset-classification process should include information on how to change data classification.

Table 35-2 summarizes the typical classification schemas for the two types of organizations discussed in this section.

Key Topic

Table 35-2 Classification Schema

Military/Government Classification	Commercial Classification	Damage Degree
Top Secret	Confidential	Grave damage
Secret	Private	Severe damage
Confidential	Sensitive	Damage
Unclassified	Public	Nonsignificant damage

Personally Identifiable Information and Protected Health Information

The United States government and many regulations require organizations to identify ***personally identifiable information (PII)*** and ***protected health information (PHI)*** and handle them in a secure manner. Unauthorized release or loss of such data could result in severe fines and penalties for the organization. Given the importance of PII and PHI, regulators and the government want to oversee the usage more efficiently.

Key Topic

PII

According to the Executive Office of the President, Office of Management and Budget (OMB) and the US Department of Commerce, Office of the Chief Information Officer, PII refers to "information which can be used to distinguish or trace an individual's identity." The following are a few examples:

- Individual's name
- Social Security number
- Biological or personal characteristics, such as an image of distinguishing features, fingerprints, X-rays, voice signature, retina scan, and the geometry of the face
- Date and place of birth
- Mother's maiden name
- Credit card numbers
- Bank account numbers
- Driver license number
- Address information, such as email addresses or street addresses, and telephone numbers for businesses or personal use

TIP The source of PII can be from many different industries. These include ***financial information*** collected by organizations such as banks and insurance companies. It can also be a result of ***government data*** that was collected by a specific government entity such as the Internal Revenue Service or Department of Defense. Of course, the most common source is ***customer data***. This data is typically collected by online retail organizations.

Key Topic

PHI

HIPAA requires health-care organizations and providers to adopt certain security regulations for protecting ***health information***. The Privacy Rule calls this information "protected health information," or PHI. This information includes, but is not limited to, the following:

- Individual's name (that is, patient's name)
- All dates directly linked to an individual, including date of birth, death, discharge, and administration
- Telephone and fax numbers
- Email addresses and geographic subdivisions such as street addresses, zip codes, and county
- Medical record numbers and health plan beneficiary numbers
- Certificate numbers or account numbers
- Social Security number
- Driver license number
- Biometric identifiers, including voice or fingerprints
- Photos of the full face or recognizable features
- Any unique number-based code or characteristic
- The individual's past, present, and future physical or mental health or condition
- The provision of health care to the individual, or the past, present, or future payment for the provision of health care to the individual

Key Topic

Privacy Enhancing Technologies

In today's world, privacy is very hard to come by. Our everyday lives are on the Internet—some more than others, some by choice and some not. Everyone has a digital footprint. Our information is out there on the Internet regardless of how hard we work to protect it. Accepting that fact is the first step in enabling yourself to address a personal privacy incident. In this section, we discuss privacy enhancing technologies.

Data minimization is a term used to explain a concept or approach to privacy design. The overall concept of data minimization is simply to minimize the amount of your personal information that is consumed by online entities. Data minimization

is a privacy tool that is used in many different ways. For instance, a website may choose to not store your personal information if it is not needed—as opposed to many that store it and even resell it for a profit. Additionally, it can be used to develop policies regarding the amount of time the data that is collected about you is actually maintained before being permanently deleted. Individuals can also use tools that will clear information from applications such as web browsers. As you know, web browsers collect a large amount of data, which could in turn be compromised. Minimizing this data reduces the risk of such a compromise.

Another privacy enhancing technology concept is ***data masking***. The goal of data masking is to protect or obfuscate sensitive data. This goal must be achieved while not rendering the data unusable in any way. An example of data masking being used in a real-world environment would be data that is displayed on terminal screens in banks or doctors' offices. Social Security numbers can be masked to show only the last four digits so that they can be used for verification purposes while not exposing the full Social Security number. A similar form of data obfuscation that is used in privacy enhancing technology is ***tokenization***. The approach of tokenization is to utilize a token to replace or obscure the data in a reversable manner.

Anonymization is one of the most critical concepts in privacy enhancing technology. The value of data to an adversary is based on how that data can be used for monetary or other tactical purposes. Being able to correlate various data obtained by adversaries can provide them with the ability to expose an individual or organization. Anonymizing the data allows for it to be used for legitimate purposes while not exposing the identity of the data owner. An example of the need for data anonymization is in the health-care industry. The idea of data anonymization is to decouple the actual sensitive data from the individual to minimize the impact of compromise. In many cases, the actual deployment of anonymization is more in line with ***pseudo-anonymization***, where the data is not completely anonymized. An example would be where the nonsensitive data is not anonymized and is able to be used for other business purposes.

NOTE Redaction is often used as a method of obfuscating data. The methods involve replacing some or all of the sensitive data for security and privacy. An example of this is the use of an asterisk when viewing a password field on an application.

Roles and Responsibilities

The previous sections described the pillars of an access control process and emphasized the importance of correctly classifying data and assets. Who decides whether a set of data should be considered confidential? Who is ultimately responsible in the case of unauthorized disclosure of such data?

Because data is handled by several people at different stages, it is important that an organization build a clear role and responsibility plan. By doing so, the organization maintains accountability and responsibility, reducing confusion and ensuring that security requirements are balanced with the achievement of business objectives.

Regardless of the user's role, one of the fundamental principles in security is that maintaining the safekeeping of information is everyone's responsibility.

Key Topic

The following key concepts are related to security roles and responsibilities:

- The definition of roles is needed to maintain clear responsibility and accountability.
- Protecting the security of information and assets is everyone's responsibility.

The following roles are commonly used within an organization, although they might be called something different, depending on the organization. Additionally, depending on the size of the organization, an individual might be assigned more than one role.

- **Executives and senior management:** They have the ultimate responsibility over the security of data and assets. They should be involved in and approve access control policies.
- ***Data owner*:** The data owner, also called the information owner, is usually part of the management team and maintains ownership of and responsibility for a specific piece or subset of data. Part of the responsibility of this role is to determine the appropriate classification of the information, ensure that the information is protected with controls, periodically review classification and access rights, and understand the risk associated with the information.
- ***Data custodian/steward*:** The data custodian/steward is the individual who performs day-to-day tasks on behalf of the data owner. This person's main responsibility is to ensure that the information is available to the end user and that security policies, standards, and guidelines are followed.
- ***Data controller*:** The data controller is the individual who has the greatest responsibility for data privacy protection. This person's main responsibility is to control how the data is used by applying specific procedures for the data processes.
- ***Data processor*:** The data processor is the individual who processes the data from the data controller. The data processer must follow the processes put in place by the data controller on how the data is to be used.
- **System owner:** The system owner is responsible for the security of the systems that handle and process information owned by different data owners.

This person's responsibility is to ensure that the data is secure while it is being processed by the system he or she owns. The system owner works closely with the data owner to determine the appropriate controls to apply to data.

- **Security administrator:** The security administrator manages the process for granting access rights to information. This includes assigning privileges, granting access, and monitoring and maintaining records of access.
- **End user:** This role is for the final users of the information. They contribute to the security of the information by adhering to the organization's security policy.

Besides these roles, several others can be seen in larger organizations, including the following:

- **Security officer:** This person is in charge of the design, implementation, management, and review of security policies and organizing and coordinating information security activities.
- ***Data protection officer (DPO)*:** The data protection officer is an organizational leadership role that is responsible for the overall protection and adherence to data protection processes within the organization. This is a role that is required by General Data Protection Regulation (GDPR) in the European Union (EU).
- **Information systems security professional:** This person is responsible for drafting policies, creating standards and guidelines related to information security, and providing guidance on new and existing threats.
- **Auditor:** This person is responsible for determining whether owners, custodians, and systems are compliant with the organization's security policies and providing independent assurance to senior management.

Key Topic Information Lifecycle

The GDPR defines the ***information lifecycle*** in four different phases:

1. Collection of data
2. Storage of data
3. How data is used
4. Disposal of data

Although the steps of the information lifecycle are sometimes named differently based on the source, the phases themselves are still the same.

The collection of data is the phase where the data is consumed by the data processor. GDPR states that when collecting data, there must be a defined consent from the data owner as well as a clear definition of how the data will be used. The overall intent of this is to follow the principle of collecting only data that is necessary and not overcollecting.

After the data is collected, of course, it must go somewhere. It must be stored or maintained in some way. This is the next phase of the information lifecycle: storing and securing the collected data. Defining things like where the data is stored and how long it should be stored are important aspects of this phase of the information lifecycle. GDPR requires that the organization collecting the data maintains the accountability of the security of that data.

Of course, if data is being collected and stored, it should be for a specific purpose. That is the next phase of the information lifecycle: how the data that was collected is actually used. Part of this phase is defining who has access to the data and what they can use it for.

The final phase of the information lifecycle is the disposal of the data. If you do not specifically define how and when the data is to be disposed of, it might linger out there forever, which increases the likelihood of the data being compromised at some point. This is a concept discussed in more detail in Chapter 15, "Understanding the Importance of Physical Security Controls."

Impact Assessment

Before data is collected, stored, secured, and disposed of throughout the information lifecycle, it is important to understand how that data, if compromised, could impact the privacy of the individuals whose data it holds. To accomplish this, an organization should complete an ***impact assessment*** on any new projects that are to be instated where data will be collected as well as any time the scope of the data use will change. These impact assessments are also sometimes called *Privacy Impact Assessments (PIA)* or *Data Privacy Impact Assessments (DPIA)*. The result of an impact assessment should produce some kind of report that will identify specific high risks to the data subjects and provide recommendations on how that risk can be minimized.

Terms of Agreement

Another privacy concept that has been adopted by the GDPR is the ***terms of agreement***. In many cases, it is called the data processing agreement. The overall purpose of the data processing agreement is to protect the personal information and the individuals the data is about. The agreement is actually a legal contract that is agreed upon by any entities that will fit the role of data processor in the information lifecycle.

Privacy Notice

Along with the agreement of how data will be collected, utilized, and processed by an organization, the organization must also provide notification to the individuals it is collecting data from or about. Again, this is a requirement for the General Data Protection Regulation. GDPR ensures that individuals are notified about how their data is being used. This is done via a ***privacy notice***. The notice itself is a document sent from the collecting organization stating how it is conforming to data privacy principles.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 35-3 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 35-3 Key Topics for Chapter 35

Key Topic Element	Description	Page Number
Section	Organizational Consequences of Privacy and Data Breaches	940
Section	Data Types and Asset Classification	941
Table 35-2	Classification Schema	942
Section	Personally Identifiable Information	943
Section	Protected Health Information	944
Section	Privacy Enhancing Technologies	944
Paragraph	Roles and responsibilities	946
Section	Information Lifecycle	947

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

intellectual property (IP) theft, reputation damage, identity theft, fines, public notifications and disclosures, escalation, confidential, proprietary, sensitive,

public, critical, personally identifiable information (PII), protected health information (PHI), financial information, government data, customer data, health information, data minimization, data masking, tokenization, anonymization, pseudo-anonymization, data owner, data custodian/steward, data controller, data processor, data protection officer (DPO), information lifecycle, impact assessment, terms of agreement, privacy notice

Review Questions

Answer the following review questions. Check your answers with the answer key in Appendix A.

1. Unauthorized access to _____ information could cause severe damage to the organization.
2. A compromise of _____ data could cause grave damage to national security?
3. Telephone and fax numbers are a form of which type of information?
4. Medical records are a form of which type of information?
5. The term _____ is used to explain reducing the amount of data as a privacy tool.
6. What form of data obfuscation is performed by replacing data in a reversable manner?
7. What is the role of the individual who has the greatest responsibility in data privacy?
8. What leadership role in an organization is responsible for the overall protection and adherence to the data protection process?

This page intentionally left blank



Final Preparation

The first 35 chapters of this book cover the technologies, protocols, design concepts, and considerations required for your preparation for passing the CompTIA Security+ SY0-601 certification exam. These chapters cover the information that is necessary to pass the exam. However, most people need more preparation than simply reading the first 35 chapters of this book. This chapter, along with the Introduction of the book, suggests hands-on activities and a study plan that will help you complete your preparation for the exam.

Hands-on Activities

The CompTIA Security+ SY0-601 certification exam is not a hands-on exam. However, one of the best ways to learn about cybersecurity vulnerabilities, threats, and techniques is to perform some hands-on exercises. You can practice using basic tools like Wireshark, tshark, tcpdump, nmap, and others. A good place to start is with Kali Linux (https://kali.org) to become familiar with some of the most common open-source attack tools and Security Onion (https://securityonion.net) to become familiar with tools related to incident response. You can also set up the learning environment called WebSploit Labs (https://websploit.org) created by Omar Santos.

TIP Although building your own test lab is beyond the scope of this book, you might want to check out http://h4cker.org/lab. It will guide you through building your own lab and using many of these tools. The most effective way to learn the skills necessary to pass the exam is to build your own lab, break it, and fix it.

Suggested Plan for Final Review and Study

This section lists a suggested study plan from the point at which you finish reading this book through Chapter 35 until you take the CompTIA Security+ SY0-601 certification exam. You can ignore this four-step plan, use it as is, or modify it to better meet your needs:

Step 1. **Review key topics:** You can use the table at the end of each chapter that lists the key topics in each chapter or just flip the pages looking for key topics.

Step 2. **Review testable content:** CompTIA maintains a list of testable content known as the CompTIA Security+ SY0-601 certification exam blueprint. Review it and make sure you are familiar with every item that is listed. You can download a copy at https://www.comptia.org/certifications.

Step 3. **Study "Review Questions" sections:** Go through the review questions at the end of each chapter to identify areas in which you need more study.

Step 4. **Use the Pearson Test Prep software to practice:** The Pearson Test Prep practice test software provides a bank of unique exam-realistic questions available only with this book.

The Introduction of this book contains the detailed instructions on how to access the Pearson Test Prep practice test software. This database of questions was created specifically for this book and is available to you either online or as an offline Windows application. As covered in the Introduction, you can choose to take the exams in one of three modes: Study Mode, Practice Exam Mode, or Flash Card Mode.

Summary

The tools and suggestions listed in this chapter have been designed with one goal in mind: to help you develop the skills required to pass the CompTIA Security+ SY0-601 certification exam and gain the skills needed to start your cybersecurity operations career. This book has been developed from the beginning both to present you with a collection of facts and to help you learn how to apply those facts. Regardless of your experience level before reading this book, it is our hope that the broad range of preparation tools, and even the structure of the book, will help you pass the exam with ease. We wish you success in your exam and hope that our paths cross again as you continue to grow in your cybersecurity career.

Glossary of Key Terms

Numerics

.p7b A certificate format that can be used by itself in IIS as the basis for S/MIME and single sign-on.

5G Fifth-generation advanced wireless network technology developed based on 802.11ac IEEE wireless standard.

802.1X An IEEE standard that defines port-based network access control (PNAC). Not to be confused with 802.11x WLAN standards, 802.1X is a data link layer authentication technology used to connect hosts to a LAN or WLAN. It allows you to apply a security control that ties physical ports to end-device MAC addresses, and prevents additional devices from being connected to the network. It is a good way of implementing port security, much better than simply setting up MAC filtering.

A

acceptable use policies Policies that define the rules restricting how a computer, network, or other system may be used.

access control list (ACL) A list of permissions attached to an object specifying what level of access a user, users, or groups have to that object. When you're dealing with firewalls, an ACL is a set of rules that apply to a list of network names, IP addresses, and port numbers.

access control model A collection of policies to determine the level of access that a subject (user or system) has on a resource (the system, application, or data to be protected). There are four major types of access control models: Mandatory Access Control (MAC), Role-Based Access Control (RBAC), Discretionary Access Control (DAC), and Rule-Based Access Control (RBAC or RB-RBAC).

access control vestibule Entries with panels built from prefabricated composite or metal that are used as a way to control ingress/egress of a building, allowing security to have visibility into persons accessing the facility. They can also be used to control the heat and airflow in their facilities.

access policies Policies that are implemented to allow or deny access to an entity. This could be a network, file, or any other kind of resource.

account audit The process of auditing the permissions that are assigned to accounts on a system. Permissions often are added when needed but never removed when not needed.

account permissions Permissions, such as file and printer access, that can be assigned to individual users or to groups.

accounting The process that logs login session statistics and usage information when a user accesses an application, system, or device.

acquisition In digital forensics, the process of collecting specific data related to an attack, intrusion, or investigation, which can include computer media and other devices that store electronic data.

Active Directory A Microsoft directory service that authenticates and authorizes users and computers.

active reconnaissance An attacker's method that is carried out on a target mostly by using network and vulnerability scanners.

active/active A load-balancing scenario in which each device performs work simultaneously, thus sharing the load.

active/passive A load-balancing scenario in which one device actively performs work while the other works in a standby mode.

address space layout randomization (ASLR) A programming method involving random arrangement of different address spaces used by a program (or process). It helps prevent the exploitation of buffer overflows, remote code execution, and memory corruption vulnerabilities. It also can aid in protecting mobile devices (and other systems) from exploits caused by memory-management problems.

administrator accounts Accounts on a system with higher-level privileges. They are similar to root accounts on a Linux system.

advanced persistent threat (APT) A sophisticated attack that can remain undetected for a long time. Also, a government (state actor) attack is often also referred to as an APT, although this term can also refer to the set of computer-attacking processes themselves. Often, an APT entity has the highest level of resources, including open-source intelligence (OSINT) and covert sources of intelligence.

adversary tactics, techniques, and procedures (TTPs) The tactics, techniques, and procedures used by attackers to compromise a system or a network.

air gap A concept that refers to the gap or lack of connection between a computer and other networks. Because the computer isn't directly connected to the network, it can't be attacked through the network.

allow list A list of allowed applications or functions that are accessible to a specific resource, such as another application, a system, or a user. The list is inclusive; if the application is not listed, access is denied.

always-on VPN A VPN client that immediately and automatically establishes a VPN connection when an Internet connection is made.

annualized loss expectancy (ALE) The total expected loss in dollars per year due to a specific incident.

annualized rate of occurrence (ARO) The number of times per year that a specific incident occurs.

anonymization A method of obfuscating data such that the data can be used for legitimate purposes while not exposing the identity of the data owner.

antimalware Software that protects against infections caused by many types of malware, including all types of viruses, as well as rootkits, ransomware, and spyware.

antivirus software A computer program used to prevent, detect, and remove malware. Also known as anti-virus software or antimalware.

anything as a service (XaaS) Cloud-based services that don't fall into SaaS, PaaS, or IaaS. For example, when a large service provider integrates its security services into the company/customer's existing infrastructure, it is often referred to as Security as a service (SECaaS).

API inspection and integration A capability available in all cloud computing environments. It allows for better automation of workflow deployment. These integrations typically need to be enabled in the environment to utilize them.

application approved list An index of approved software applications or executable files that are permitted to be present and active on a computer system.

application block list/deny list An index or list of undesirable or unauthorized programs used to prevent their execution.

application DDoS attacks Attacks that target the resources of Layer 7 applications and often leverage known vulnerabilities against specific software.

application programming interface (API) A programmatic framework that enables other systems to interact with an application; however, lack of adequate controls and monitoring make effective security testing of APIs difficult to automate, which makes them vulnerable targets.

application provisioning A process to adequately and securely deploy an application on-premises or in the cloud.

application scanners Devices used to assess application-specific vulnerabilities and operate at the upper layers of the OSI model.

Arduino An open-source electronics platform based on easy-to-use hardware and software.

arp Command that displays and modifies entries in the Address Resolution Protocol (ARP) cache, which contains one or more tables that are used to store IP addresses and their resolved Ethernet physical addresses.

ARP cache poisoning When an attacker manipulates the ARP cache on a host to redirect traffic and perform an on-path attack.

artifacts Remnants of an intrusion that can be identified on a host or network.

asset management A policy for onboarding and offboarding devices; it specifies how they are registered and activated and how they are later decommissioned.

asset value The monetary value of an asset.

asymmetric encryption A process that uses a public-key and private-key pair to encrypt and decrypt messages when communicating.

ATT&CK A set of matrices created by MITRE to document and explain the adversarial tactics and techniques used by attackers to compromise systems and networks.

attestation A process that serves to bear witness and to confirm, authenticate, verify, and document.

attribute-based access control (ABAC) An access model that is dynamic and context-aware. Access rights are granted to users through the use of multiple policies that can combine various user, group, and resource attributes together.

attributes Characteristics that authenticate a user in either a physical or behavioral manner.

audio steganography A technique used to transmit hidden information by modifying an audio signal in an imperceptible manner.

authenticated mode An encryption type that includes authentication (or authenticated encryption [AE]) and authenticated encryption with associated data (AEAD). With AE, the requirement is to both protect the privacy of the message and to ensure authenticity. A method for achieving both of these goals at the same time is called authenticated encryption.

authentication The process or action of proving something to be true or valid, verifying the identity of a user or process.

authentication application A program that generates security codes for signing into assets.

Authentication Header (AH) As specified in RFC 4302, a protocol that defines an optional packet header to be used to guarantee connectionless integrity and data origin authentication for IP packets and to protect against replays.

authorization The action of approving or being authorized to determine access levels.

authorized hacker A nonmalicious hacker—for example, an IT person who attempts to "hack" into a computer system before it goes live to test the system.

automated courses of action A DevOps environment component for secure provisioning and deprovisioning of software, services, and infrastructure.

automated indicator sharing (AIS) An automated way to share indicators of compromise (IOCs) and threat intelligence information.

automation The technology and processes of executing a task without human intervention.

Autopsy A digital forensics platform and graphical interface to The Sleuth Kit and other digital forensics tools. It is used by law enforcement, military, and corporate examiners to investigate what happened on a computer.

B

backdoor A method used in computer programs to bypass normal authentication and other security mechanisms in place.

background checks Procedures typically done when an employee is going through the hiring process. Background checks validate a candidate's criminal record and financial situation, for example.

bandwidth monitor Software that tracks bandwidth use over all areas of the network, devices, applications, servers, WAN, and Internet links.

barricade A defensive barrier constructed to stop or obstruct passage.

base group A Common Vulnerability Scoring System (CVSS) group that represents the intrinsic characteristics of a vulnerability that are constant over time and do not depend on a user-specific environment.

baseband radio The original frequency range of transmission signal before it is modulated. It can also refer to the type of data transmission in which analog data is sent over a single nonmultiplex channel.

baseline configuration (baselining) A method used to assess the current security state of computers, servers, network devices, and the network in general after a minimum desired state of security is defined.

baseline reporting The process of reporting the security state of computers, servers, network devices, and the network after a baseline has been determined.

baselining See *baseline reporting*.

Bash A Linux/UNIX-based scripting shell and framework.

binary A compiled program that can be read and executed by a computer.

biometric attribute The element (or attribute) to perform biometric-based authentication such as face recognition, retina scan, or fingerprint scan.

biometrics Security controls that provide a unique way of making sure that people are who they say they are by monitoring/matching human characteristics such as a fingerprint, retina, or voice.

birthday attack An attack on a hashing system that attempts to send two different messages with the same hash function, causing a collision.

black-box testing A way of testing the internal workings of an application or system where the tester has no knowledge of the system being tested.

block cipher An encryption method that applies a deterministic algorithm along with a symmetric key to encrypt a block of text instead of encrypting one bit at a time as in stream ciphers.

block list/deny list A list used to deny individual application access—a common method used when working with email, and by antivirus and HIDS programs.

blockchain A specific type of database. It differs from a typical database in the way it stores information; a blockchain stores data in blocks that are then chained together. As new data comes in, it is entered into a fresh block. Once the block is filled with data, it is chained onto the previous block, which makes the data chained together in chronological order.

blue team A term used to identity the defenders of an organization. Blue teams typically include the computer security incident response team (CSIRT) and information security (InfoSec) team.

bluejacking Sending unsolicited messages to Bluetooth-enabled devices such as mobile phones and tablets.

bluesnarfing Accessing information without authorization from a wireless device through a Bluetooth connection.

bollard A standalone post used for physical security purposes. It is typically steel, short, and sturdy, and anchored in a hard surface such as concrete.

boot attestation A process that allows a remote platform to measure and report its system state in a secure way to a third party.

boot integrity The reliability of the operating system and loading mechanism during the booting process; it can be checked using a secure method.

botnet A large group of compromised systems known as robots or simply bots.

bots Compromised computers (also known as zombies) that are part of a larger group called a botnet. They are used to distribute malware across the Internet.

Bridge Protocol Data Unit (BPDU) guard A mechanism used to prevent receipt of BPDUs from a connected device such as a client reflecting BPDUs back to the switch.

bring your own device (BYOD) A device policy where employees bring a device of their choosing for both work and personal purposes.

brute-force attack A password attack where every possible password is attempted.

buffer overflow A situation that occurs when a process stores data outside the memory that the developer intended.

bug bounties The recognition and compensation provided by an organization to security researchers for reporting security vulnerabilities (which are basically bugs in code or hardware).

business continuity plan (BCP) A current, tested plan in the hands of all personnel responsible for carrying out any part of that plan for the purpose of giving your organization the best shot at success during a disaster.

business partners In the context of this book and the Security+ blueprint, the specific concerns of business partners when addressing third-party risk management.

business partnership agreement (BPA) A type of contract that can establish the profits each partner will get, what responsibilities each partner will have, and exit strategies for partners.

C

cable locks Devices typically deployed to secure equipment to stable heavy equipment in a room that is not easily removed. For example, you would use a computer cable lock that connects a laptop computer to a desk or a projector to its cart or stand.

California SB 1386 Enacted in 2003, a law that requires California businesses that store computerized personal information to immediately disclose breaches of security.

camouflage The complicating of source code to make it more difficult for people to understand. See also *obfuscation*.

captive portal A method used by hotels, coffee shops, etc., that directs users to a web page for authentication (typically through email) prior to normal Internet use. The whole point of the technology is to be able to track users who access the free wireless network. If the user performs any suspect actions, that user can be traced by way of email address, IP address, and MAC address, in addition to other means if multifactor authentication is used.

capture the flag A method of user awareness training where students play in a red team/blue team scenario.

card cloning attack An attack method where the attacker clones a credit card, a smartphone SIM card, or a building access badge or card.

cat A Linux command that copies standard input to standard output.

CCMP Counter Mode with Cipher Block Chaining Message Authentication Code (CBC-MAC) Protocol; an encryption protocol used with WPA2 that addresses the vulnerabilities of TKIP and meets the requirements of IEEE 802.11i.

Center for Internet Security A nonprofit organization that was established in 2000 with the overall goal to provide security best practice guidance for enhancing the security of cyberspace.

certificate authority (CA) An entity (usually a server) that issues certificates to users.

certificate revocation list (CRL) A list of digitally signed certificates revoked by the certificate authority for security purposes. If a certificate is compromised, it is revoked and placed on the CRL. CRLs are later generated and published periodically.

certificates Digitally signed electronic documents that bind a public key with a user identity.

chain of custody A process that provides assurances that evidence has been controlled and handled properly after collection.

Challenge-Handshake Authentication Protocol (CHAP) An authentication scheme used by the Point-to-Point Protocol (PPP), which in turn is the standard for dial-up connections. It uses a challenge-response mechanism with one-way encryption.

change control The process that is put in place to handle requests to make changes to a system in a more efficient and coordinated manner.

change management A structured way of changing the state of a computer system, network, or IT procedure.

chmod The Linux command and system call that is used to change the access permissions of file system objects.

choose your own device (CYOD) A device policy where employees select a device from a company-approved list.

cipher suite A set of algorithms that help secure a network connection that uses Transport Layer Security (TLS). The set of algorithms that cipher suites usually contain include a key exchange algorithm, bulk encryption algorithm, and message authentication code (MAC) algorithm.

cleanup The process of completely removing any residual files or data from target systems after the testing phases of a penetration testing engagement are complete.

client-side execution Anything that is being performed (a command, script, or otherwise) at the client end of the communication. Typically executed on the client's browser rather than on the web server, it allows for more responsive web applications.

client-side validation The ability to properly handle application and user input to prevent a security vulnerability and client-side execution.

closed circuit television (CCTV) A video surveillance system that uses video cameras to monitor specific areas.

cloud access security broker (CASB) A tool that is utilized in organizations to control access to and use of cloud-based computing environments.

cloud backups Duplicate copies of your computer or entire set of computers saved to cloud-based storage services.

cloud computing A way of offering on-demand services that extend the capabilities of a person's computer or an organization's network. These might be free services, such as personal browser-based email from various providers, or they could be offered on a pay-per-use basis, such as services that offer data access, data storage, infrastructure, and online gaming. A network connection of some sort is required to make the connection to the cloud and gain access to these services in real time.

Cloud Controls Matrix A framework established by the Cloud Security Alliance for cloud computing.

Cloud Security Alliance CSA; a nonprofit organization established in 2008 with the goal of promoting best security best practices in cloud computing environments.

code reuse The act of reusing third-party, open-source software, or code developed internally by an organization.

code signing Digitally signing executables and scripts to confirm the software author and guarantee that the code has not been altered or corrupted since it was signed.

cold aisle In a hot aisle/cold aisle data center design, the rows are composed of rack fronts.

cold site A location belonging to an organization that has tables, chairs, bathrooms, and possibly some technical setup—for example, basic phone, data, and electric lines. Otherwise, a lot of configuration of computers and data restoration is necessary before the site can be properly utilized. This type of site is used only if a company can handle the stress of being nonproductive for a week or more.

collision A situation that occurs when two different files end up using the same hash, which is possible with less secure hashing algorithms.

command and control The controlling master computer directing the actions of a botnet, which distributes Internet malware.

Common Security Advisory Framework (CSAF) A standard that enables different stakeholders across different organizations to share critical security-related information in a single format, speeding up information exchange and digestion. CSAF is also the name of the technical committee in the OASIS standards organization. CSAF is the successor of the Common Vulnerability Reporting Framework (CVRF). You can obtain additional information about CSAF at csaf.io.

Common Vulnerabilities and Exposures (CVE) A standard created by MITRE (www.mitre.org) that provides a mechanism to assign an identifier to vulnerabilities so that you can correlate the reports of those vulnerabilities among sites, tools, and feeds.

community cloud A mix of public and private cloud deployments where multiple organizations can share the public portion.

compensating controls Mechanisms put in place to satisfy security requirements that are either impractical or too difficult to implement. For example, instead of using expensive hardware-based encryption modules, an organization might opt to use network access control (NAC), data loss prevention (DLP), and other security methods. Or, on the personnel side, instead of implementing segregation of duties, an organization might opt to do additional logging and auditing. Also known as alternative controls.

compiler A program that translates, verifies, and processes source code created in a specific programming language. The compiler converts the programming language instructions (source code) into a machine-code form in order for it to be read and executed by the system.

compile-time errors Program errors that occur while the program is being compiled.

conditional access An access control model where access is granted based on specific criteria requirements.

confidential A classification of information where unauthorized access to the information would cause damage to national security.

configuration management An ongoing process created with the goal of maintaining computer systems, servers, network infrastructure, and software in a desired, consistent state.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

configuration reviews The process of reviewing system configurations to reveal potential security problems.

container security The use of various tools for securing containerized cloud computing environments. Some of these tools are native to the actual cloud computing environment, and some are third-party solutions and run on those environments.

containers Logical units of software that package applications and all the dependencies needed to run it. Containers are lightweight, standalone executable software packages that include code, runtime environments, system tools, and related software libraries. Two of the most popular examples of container solutions are Docker and Linux LXC containers.

containment The methodology whereby access to information, files, systems or networks is controlled.

content URL/filtering Using a program to screen and/or exclude access to web pages, URLs, or email deemed objectionable.

continuity of operations planning (COOP) A federal initiative to encourage people and departments to plan to address how critical operations will continue under a broad range of circumstances.

continuous delivery A software development process in which developers produce software in short cycles while making sure that the software is reliable and secure.

continuous deployment The automation of the application deployment, provisioning, and underlying network components and infrastructure.

continuous integration A software development practice in which programmers merge, test, and deploy code changes in a central repository multiple times a day or several times per week.

continuous monitoring A DevOps environment component that ensures applications and systems are operating correctly and securely.

continuous validation A DevOps environment component where applications and code must be validated in an automated fashion.

copy backups The methodology of making a copy of a set of backups.

corporate-owned, personally enabled (COPE) A device policy in which the company supplies employees with a phone that can also be used for personal activities.

corrective controls Controls used after an event. They limit the extent of damage and help the company recover from damage quickly. Tape backup, hot sites, and other fault tolerance and disaster recovery methods are also included here. They are sometimes referred to as compensating controls.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

counter mode An encryption mode that uses an arbitrary number (the counter) that changes with each block of text encrypted. The counter is encrypted with the cipher, and the result is XOR'd (exclusive OR'd) into ciphertext. Because the counter changes for each block, the problem of repeating ciphertext that results from the Electronic Code Book method is avoided.

counterintelligence Information gathered and activities conducted to protect against espionage, other intelligence activities, or sabotage conducted by or on behalf of other elements.

credential harvesting The attacking technique or activities of grabbing legitimate usernames and even passwords to gain access to systems to steal information or to use them for malicious purposes.

criminal syndicates Centralized enterprises run by people motivated mainly by money (organized crime).

critical A description of data that is important to the continued function of a business. Loss of this type of data would result in significant monetary loss.

Crossover Error Rate (CER) The point where the false rejection rate (FRR) and the false acceptance rate (FAR) are equal.

cross-site request forgery (XSRF) A type of vulnerability where an attacker lures the targeted user to execute unwanted actions on a web application. Threat-performing XSRF attacks leverage the trust that the application has in the targeted user.

cross-site scripting (XSS) A web application vulnerability where an attacker can redirect a user to a malicious site, steal session cookies, or steal other sensitive information.

cryptographic attacks Attacks against cryptographic implementations or against crypto algorithms.

cryptomalware An advanced form of ransomware that leverages advanced encryption techniques to prevent files from being decrypted without a unique key.

Cuckoo Open-source software/sandbox for automating analysis of suspicious files.

curl A Linux command-line tool to transfer data to or from a server, using any of the supported protocols: HTTP, FTP, IMAP, POP3, SCP, SFTP, SMTP, TFTP, TELNET, LDAP, or FILE.

customer data Personally identifiable information that was collected by online or in-person retail organizations.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

CVSS Common Vulnerability Scoring System; an industry standard used to convey information about the severity of vulnerabilities.

cyber kill chain A series of steps that trace stages of a cyber attack from the early reconnaissance stages to the exfiltration of data. The kill chain allows you to understand and combat ransomware, security breaches, and advanced persistent threats (APTs).

cybersecurity insurance A type of risk transference (also known as risk sharing) that an organization can purchase to protect, for example, a group of servers in a data center.

D

dark web A subset of the deep web where many threat actors perform malicious activities, such as selling stolen credit card numbers, health records, and other personal information.

data at rest Inactive data that is archived—backed up or stored in cloud storage services.

data controller The individual who has the greatest responsibility for data privacy protection. This person's main responsibility is to control how the data is used by applying specific procedures for the data processes.

data custodian/steward The individual who performs day-to-day tasks on behalf of the data owner. This person's main responsibility is to ensure that the information is available to the end user and that security policies, standards, and guidelines are followed.

data exposure The unauthorized or unnecessary exposure of sensitive information.

data in transit/motion Data that crosses the network or data that currently resides in computer memory. Also known as data in motion.

data in use/processing Actively used data undergoing constant change; for example, it could be stored in databases or spreadsheets.

data input The process of entering data into a system or an application.

data loss prevention (DLP) A system that performs content inspection and is designed to prevent unauthorized use of data as well as prevent the leakage of data outside the computer (or network) in which it resides.

data masking A privacy enhancing technology designed to protect or obfuscate sensitive data.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

data minimization A method of minimizing the amount of personal information that is consumed by online entities.

data owner Also called the information owner; a person who is usually part of the management team and maintains ownership of and responsibility over a specific piece or subset of data.

data processor The individual who processes data from the data controller. This person must follow the processes put in place by the data controller on how the data is to be used.

data protection The act of monitoring data in processing, data in transit/motion, and data at rest.

data protection officer (DPO) The organizational leadership role that is responsible for the overall protection and adherence to data protection processes within the organization.

data recovery The process of restoring lost data, such as restoring a corrupt file from a backup. In the context of forensics, data recovery goes further. Even without backups, it's often possible to recover data that has been intentionally or accidentally deleted.

data retention A policy that states how long data must be stored by an organization.

data sanitization The process of irreversibly removing or destroying data stored on a memory device (hard drives, flash memory/SSDs, mobile devices, CDs, DVDs, and so on) or in hard copy form.

data sovereignty A concept that refers to any information (data) that has been converted and stored in a digital form.

data storage The process of storing information in a system or digital media.

dd A command-line utility for Linux operating systems whose primary purpose is to convert and copy files.

dead code Software that is no longer in use, but the source code or binary for it still exists in the system and has not been removed appropriately.

degaussing A data destruction method involving the reduction or elimination of a magnetic field (or data) stored on tape and disk media such as computer and laptop hard drives, diskettes, reels, cassettes, and cartridge tapes.

deny list A list that is generally utilized at the end of an allowed list; if all items on the list are allowed, then anything not on the list is denied. The deny list, while similar to the block list, is typically used in access control systems.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

deprovision To adequately and securely remove, decommission, and purge an application and related data on-premises or in the cloud.

detective controls Controls aimed at monitoring and detecting any unauthorized behavior or hazard. These types of controls are generally used to alert to a failure in other types of controls such as preventive, deterrent, and compensating controls.

deterrent controls A control that is used by an organization to try to deter a threat actor from executing an offensive assault on its environment. Deterrent controls are similar to preventive controls in the sense that the primary objective is to prevent an incident from occurring. Unlike preventive controls, however, the rationale behind deterrent controls is to discourage attackers from proceeding just because a control is in place.

development The environment where you create code on your computer or in the cloud.

DHCP snooping A security feature that protects against Layer 2 attacks such as DHCP spoofing and abuse.

diagrams Part of the documentation used to ensure good configuration management processes.

Diamond Model of Intrusion Analysis A cybersecurity/threat intelligence model used to analyze and track the characteristics of cyber intrusions by advanced threat actors that emphasizes the relationships and characteristics of the adversary, capabilities, infrastructure, and victims.

dictionary-based attack A password attack that uses a prearranged list of likely words, trying each of them one at a time.

differential backups Data backups that preserve data, saving only the difference in the data since the last full backup.

dig A Linux tool for querying DNS nameservers for information about host addresses, mail exchanges, nameservers, and related information.

digital rights management (DRM) The name given to a set of access control technologies that are used to control the use of proprietary hardware, software, and copyrighted works. DRM solutions are used to restrict the use, modification, and distribution of copyrighted works and the underlying systems used to enforce such policies.

digital signatures Mathematical schemes for verifying the authenticity of digital messages or documents.

directory services Customizable information stores that provide a single point from which users can locate resources and services.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

directory traversal A method of accessing unauthorized parent (or worse, root) directories. It is often used on web servers that have PHP files and are Linux- or UNIX-based, but it can also be perpetrated on Microsoft operating systems. This attack is designed to get access to files such as ones that contain passwords. Also known as path traversal or the ../ (dot-dot-slash) attack.

disaster recovery plan (DRP) A formal document created by organizations that contains detailed instructions on how to respond to unplanned incidents such as natural disasters, power outages, cyber attacks, or other disruptive events. A DRP should include information regarding redundancy, such as sites and backup, but should not include information that deals with the day-to-day operations of an organization, such as updating computers, patch management, monitoring and audits, and so on. It is important to include only what is necessary in a disaster recovery plan.

discretionary access control (DAC) An access control policy generally determined by the owner. Objects such as files and printers can be created and accessed by the owner. Also, the owner decides which users are allowed to have access to the objects, and what level of access they may have. The levels of access, or permissions, are stored in access control lists (ACLs).

disk backups A data backup and recovery method that backs up data to a hard disk storage unit.

disk encryption A technology that protects information by converting it into unreadable code that cannot be deciphered easily by unauthorized people.

disk redundancy The principle behind writing data to two or more disks at the same time.

Distinguished Encoding Rules (DER) A restricted variant of Basic Encoding Rules (BER) that allows for only one type of encoding, and has restrictive rules for length, character strings, and how elements are sorted. It is widely used for X.509 certificates. For example, certificate enrollment in Windows Servers uses DER exclusively.

distributed denial-of-service (DDoS) attack An attack in which a group of compromised systems attacks a single target, causing a denial of service to occur at that host.

diversity Adequate distance between primary and secondary (or backup) sites; this is an important disaster recovery term.

DLL injection The process of running code within the address space of another process by forcing it to load a dynamic link library (DLL). Ultimately, this can influence the behavior of a program in a way that was not originally intended.

DNS Domain Name System; a system utilized on networks to translate domain names to IP addresses.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

DNS amplification attack An attack that generates a high volume of packets ultimately intended to flood a target website.

DNS poisoning The modification of name resolution information that should be in a DNS server's cache in order to redirect client computers to incorrect websites.

DNS Security Extensions (DNSSEC) A suite of specifications that provide secure answer validation. It does this through public key cryptography. It is backward-compatible and can be deployed side by side with traditional DNS.

DNS sinkhole A deception and disruption technique used when you configure one or more DNS servers to provide false results to attackers and redirect them to areas in the network where you can observe their tactics and techniques. DNS sinkholes have been used to contain different types of malware such as the famous WannaCry ransomware and to disrupt certain malicious DNS operations in denial-of-service (DoS) and other attacks. Adversaries have also used similar techniques to perform DNS poisoning attacks to redirect systems and users to malicious destinations.

dnsenum A command-line tool that automatically identifies basic DNS records such as MX, mail exchange servers, NS, domain name servers, or the address record for a domain.

domain hijacking An attack in which the attacker changes the registration of a domain name without the permission of the original owner or registrant.

Domain Keys Identified Mail (DKIM) A method that provides gateway-based cryptographic signing of outgoing messages. It allows you to embed verification data in an email header and for email recipients to verify the integrity of the email messages.

domain name kiting The process of deleting a domain name during the five-day grace period (known as the add grace period, or AGP) and immediately reregistering it for another five-day period. This process is repeated any number of times with the end result of having the domain registered without ever actually paying for it. It is a malicious attack on the entire Domain Name System (DNS) by misusing the domain-tasting grace period. The result is that a legitimate company or organization often cannot secure the domain name of its choice.

domain reputation A technique to validate the authenticity of a domain and the services using such domains (including websites and email messages). Domain reputation is used to track known malicious domains that point to websites hosting malware or those that are used for spam, phishing, spear phishing, and other malicious activities.

domain validation (DV) The process the certificate authority uses to check the rights of the applicant to use a specific domain name.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Domain-based Message Authentication, Reporting & Conformance (DMARC) A standard that was designed to thwart spammers from spoofing your domain to send email. Spammers can counterfeit the "From" address on an email message for it to appear to come from a user in your domain. DMARC is designed to block these emails before they appear in your email inbox. DMARC also provides visibility and reports into who is sending email on behalf of your domain to make sure that only legitimate emails are received.

downgrade attack A type of attack in which a protocol (such as TLS or SSL) is downgraded from the current version to a previous version, exploiting backward compatibility.

drones An unmanned air device; they can be aerial or underwater vehicles that use embedded compute platforms.

dual supply Power that is supplied to the building via multiple paths; it ensures a single path failure does not interrupt power to the building.

dumpster diving Literally scavenging for private information in garbage and recycling containers.

dynamic code analysis The process of trying to understand the source code of a program to adequately build a series of correct inputs for test coverage. Analysis software has the capability to find security issues caused by the code's interactions with other system components.

Dynamic Host Configuration Protocol (DHCP) snooping A method to determine if an ARP packet is valid based on IP-to-MAC address bindings stored in a trusted database.

dynamic resource allocation The capability of a cloud computing environment to efficiently allocate resources to tenants based on demand. Without this capability, a cloud-based computing environment would not be feasible.

E

east-west traffic A concept referring to network traffic flow within a data center between servers. North-south refers to data transfers between the data center and outside of the network.

edge computing An ecosystem of resources and applications in new network services (including 5G and IoT).

E-discovery The process of identifying, preserving, collecting, processing, reviewing, and analyzing electronically stored information (ESI) in litigation.

elasticity The capability of an underlying infrastructure to react to a sudden increase in demand by provisioning more resources in an automated way.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

electronic locks Locks that usually have a magnetic strike plate that are energized to keep the magnet engaged until authorization is given to disengage, thereby releasing the magnetic hold.

eliciting information The act of gaining knowledge or information from people.

elliptic-curve cryptography (ECC) An approach to public-key cryptography based on the algebraic structure of elliptic curves over finite fields.

embedded systems Microprocessor-based computer hardware systems with software that is designed to perform a dedicated function, either as an independent system or as a part of a larger system.

Encapsulating Security Payload (ESP) Specified in RFC 4303; an optional packet header that can be used to provide confidentiality through encryption of the packet, as well as integrity protection, data origin authentication, access control, and optional protection against replays or traffic analysis.

encryption The process used to protect data by encoding plaintext data using cryptographic algorithms. You can encrypt data at rest, in transit, and in process.

end of life (EOL) The date when a product or service will no longer be sold or supported by a third party.

end of service life (EOSL) Typically, the last day of service for a product, meaning a third party is no longer providing service if there is a failure.

endpoint detection and response (EDR) An integrated endpoint security solution that combines real-time continuous monitoring and collection of endpoint data with rules-based automated response and analysis capabilities.

endpoint protection Often used interchangeably with endpoint security; security solutions that address endpoint device security issues, securing and protecting endpoints against zero-day exploits, attacks, and inadvertent data leakage resulting from human error.

entropy In cybersecurity, a measure of the randomness or diversity of a data-generating function.

environmental disaster Disaster caused by environmental factors, such as floods, earthquakes, lightning, or hurricanes.

environmental group A Common Vulnerability Scoring System (CVSS) group that represents the characteristics of a vulnerability, taking into account the organization's environment.

ephemeral keys Cryptographic keys that can be used more than once within a single session, such as for broadcast applications, where the sender generates only

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

one ephemeral key pair per message, and the private key is combined separately with each recipient's public key.

error handling The ability to handle specific application or program errors and conditions.

escalation The process of engaging senior analysts and other stakeholders during the incident response process.

ethical hacker An expert at breaking into systems who can attack systems on behalf of the system's owner and with the owner's consent.

event logs Records that contain a lot of useful information about the system and its users. Many analysts rely on Windows event logs to help gain context of attacker activity on a system.

evil twin A rogue and unauthorized wireless access point that uses the same service set identifier (SSID) name as a nearby wireless network, often a public hotspot.

expiration The period of time that the certificate should be considered valid; a certificate is issued with a valid from and valid to date. If the certificate is encountered any time before or after those dates, the application that is validating the certificate such as a web browser should produce an error message stating that the certificate is not valid.

exploitation frameworks Software packages that contain reliable exploit modules and other hacker technique tools such as agents used for successful repositioning.

extended validation (EV) Certificates that conduct a thorough vetting of an organization. Issuance of these certificates is strictly defined.

Extensible Authentication Protocol (EAP) Specified in IETF RFC 3748 [18]; a framework for access authentication, which supports different authentication methods that are specified as EAP methods. As described in RFC 4017 [19], it is desirable for EAP methods used for wireless LAN to support mutual authentication and key derivation.

Extensible Authentication Protocol - Flexible Authentication via Secure Tunneling (EAP-FAST) A type of Extensible Authentication Protocol authentication that uses a protected access credential instead of a certificate to achieve mutual authentication. FAST stands for Flexible Authentication via Secure Tunneling.

Extensible Authentication Protocol - Transport Layer Security (EAP-TLS) A type of Extensible Authentication Protocol authentication that uses Transport Layer Security, which is a certificate-based system that does enable mutual authentication. It does not work well in enterprise scenarios because certificates must be configured or managed on the client side and server side.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Extensible Authentication Protocol - Tunneled Transport Layer Security (EAP-TTLS) A type of Extensible Authentication Protocol authentication that uses Tunneled Transport Layer Security and is basically the same as TLS except that it is done through an encrypted channel and requires only server-side certificates.

external actor An attacker who is external to the organization (not an employee or contractor).

external risk The risk to an organization from outside an environment—for instance, that of an external threat actor.

F

facility automation A full-service commercial and industrial automatic control system for buildings, factories, and larger facilities.

fake telemetry System decoys and breadcrumbs designed to lure and trick attackers.

false acceptance rate (FAR) The process of measuring the likelihood and probability that a biometric system will authorize a person who was not meant to be authorized.

false negative A network intrusion device's inability to detect true security events under certain circumstances—in other words, a malicious activity that is not detected by the security device.

false positive A situation in which a security device triggers an alarm, but no malicious activity or actual attack is taking place. In other words, false positives are false alarms, and they are also called benign triggers. False positives are problematic because by triggering unjustified alerts, they diminish the value and urgency of real alerts. Having too many false positives to investigate becomes an operational nightmare, and you most definitely will overlook real security events.

false rejection rate (FRR) The process of measuring the likelihood and probability that a biometric system fails to recognize an authorized user.

Faraday cage An enclosure designed to block any RF signals from entering or leaving, or having effect on devices inside the cage.

Field-Programmable Gate Array (FPGA) An array of programmable logic blocks and a hierarchy of "reconfigurable interconnects" that allow the blocks to be "wired together." Logic blocks can be configured to perform complex combinational functions, or merely simple logic gates such as the AND gate, OR gate, and NOT gate.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

file and code repositories Systems used to store and track changes in source code and files.

File Transfer Protocol, Secure (FTPS) A name used to encompass a number of ways in which FTP software can perform secure file transfers. Also known as FTP over TLS, FTP over SSL, FTP/SSL, FTP-SSL, FTP-ES, and FTP Secure.

fileless malware A form of malware that functions without putting malicious executables within the file system and instead works in a memory-based environment.

financial information Personally identifiable information that can be from many different industries. It includes financial information collected by organizations such as banks, insurance companies, the government, and online retail organizations.

fines Charges imposed on an organization if it does not comply with regulations.

fire suppression The act of extinguishing fires through the application of a substance.

firmware over-the-air (OTA) updates The remote configuration and deployment of mobile devices. OTA historically refers to the deployment and configuration performed via a messaging service, such as Short Message Service (SMS), Multimedia Messaging Service (MMS), Rich Communication Service (RCS), or Wireless Application Protocol (WAP).

fog computing The decentralization of computing infrastructure by "bringing the cloud to the ground" (thus the term *fog*). This architecture enables components of the edge computing concept to easily push compute power away from the public cloud to improve scalability and performance. Fog computing accomplishes this by putting data, compute, storage, and applications near the end user or IoT device.

footprinting A by-product of attacker reconnaissance on an application, system, or network in order to find vulnerabilities that could potentially be exploited.

forensics A process that deals with the recovery and investigation of material found in digital devices.

forward proxy A proxy server that clients looking for websites, or files via an FTP connection, pass their requests through to the proxy.

FTK Imager A data preview and imaging tool that lets you quickly assess electronic evidence to determine whether further analysis with a forensic tool such as AccessData Forensic Toolkit (FTK) is warranted.

full backups The process of creating one or more full copies of all data in a specific set/system.

full tunnel A type of configuration in which all traffic is sent through the VPN tunnel back to the head end and out through the corporate network.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

full-disk encryption (FDE) The process of encrypting data as it is written to the disk and decrypting data as it is read off the disk. It is most applicable to laptops.

fuzzing An automated software testing technique that involves providing invalid, unexpected, or random data as inputs to a computer program.

G

gamification A method of user-awareness training.

General Data Protection Regulation (GDPR) A European Union (EU) law that was enacted in 2018 with an overall focus on data protection and privacy for individuals.

generator A device that provides power to spaces and other devices during complete power loss, blackouts, or in areas where standard electrical service isn't available.

geofencing The use of a virtual fence defining the boundaries of an actual geographical area.

geographical considerations Considerations that affect the laws and regulations that have been created to address data privacy and digital rights management (DRM).

geographical dispersal The process of placing compute assets in strategic locations to ensure the ability to recover in case of an attack or natural disaster.

Global Positioning System (GPS) A system that offers the ability to obtain positioning, navigation, and timing (PNT) services.

government data Personally identifiable information that was collected by a specific government entity such as the IRS or DoD.

Gramm-Leach-Bliley Act Law enacted in 1999 that enables commercial banks, investment banks, securities firms, and insurance companies to consolidate. It protects against pretexting. Individuals need proper authority to gain access to nonpublic information such as Social Security numbers.

grep A Linux command for finding matching patterns, to search for a string of characters in a specified file.

guards Persons typically deployed at ingress/egress points with the goal of reducing tailgating, verifying identity of scanned ID cards, and ensuring property does not leave the premises.

guest accounts Accounts that typically require very little access rights because they are considered a security risk to most organizations.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

H

hacktivist A combination of the terms *hack* and *activist*. As with the term *hacker*, the name *hacktivist* is often applied to different kinds of activities—from hacking for social change, to hacking to promote political agendas, to full-blown cyberterrorism.

hardening The act of applying levels of security to protect applications from intellectual property theft, misuse, vulnerability exploitation, tampering, or even repackaging by people with ill intentions. Also known as application shielding.

hardware root of trust The foundation on which all secure operations of a computing system depend.

hardware security modules (HSMs) Physical devices that act as secure cryptoprocessors. This means that they are used for encryption during secure login/authentication processes, during digital signings of data, and for payment security systems. The beauty of a hardware-based encryption device such as an HSM (or a Trusted Platform Module, or TPM) is that it is faster than software encryption.

hash A cryptographic function that is a mathematical algorithm used to map data of arbitrary size to a bit array of a fixed size.

hashing A one-way function where data is mapped to a fixed-length value. Hashing is primarily used for authentication.

head A Linux-centric command that reads the first 10 lines of any given filename.

health information Information about medical patients (health records).

Health Insurance Portability and Accountability Act Law enacted in 1996 that governs the disclosure and protection of health information.

heat maps Components of a wireless site survey that shows all wireless activity in an area.

heating, ventilating, and air conditioning (HVAC) The technology of indoor, building, and data center environmental comfort. Its goal is to provide thermal comfort and acceptable indoor air quality.

Help America Vote Act of 2002 A law that governs the security, confidentiality, and integrity of personal information collected, stored, or otherwise used by various electronic and computer-based voting systems. Its main goal was to replace punch card and lever-based voting systems.

high availability A characteristic of a system which aims to ensure an agreed level of operational performance, usually uptime, for a higher than normal period.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

high availability across zones The components of a high availability environment deployed across multiple zones to greatly reduce the risk of an outage. In cloud computing environments, high availability is addressed using the concept of regions or zones.

HMAC-based one-time password (HOTP) Hash-based password algorithm that is used as a one-time password.

hoax An attempt to deceive people into believing something that is false.

homomorphic encryption A form of encryption enabling you to perform calculations on encrypted data without decrypting it first.

honeyfiles Files used as bait intended to lure adversaries to access and then send alarms to security analysts for detection. They can also be used to potentially learn the tactics and techniques used by attackers.

honeynet One or more computers, servers, or an area of a network that does not house any important company information and is designed to lure attackers so that you can study what tools and techniques they are using in order to discover potential network vulnerabilities.

honeypot A computer, virtual machine (VM), or container that is used to attract attacker traffic to learn the adversary's tactics, techniques, and procedures. In contrast, a honeynet is one or more computers, servers, or an area of a network; a honeynet is used when a single honeypot is not sufficient. Either way, the individual computer, or group of servers, usually does not house any important company information. Various analysis tools are implemented to study attackers; these tools, along with a centralized group of honeypots (or a honeynet), are known collectively as a honeyfarm.

host-based firewall A firewall installed on each individual desktop, laptop computer, or server that controls incoming and outgoing network traffic and determines whether to allow it into a particular device.

host-based intrusion detection system (HIDS) An application that operates on information collected from individual computer systems. It can detect and alert on malicious activity but cannot stop this activity.

host-based intrusion prevention system (HIPS) A system that is capable of monitoring and analyzing the internals of a computing system "server" as well as the network packets on its network. It can prevent malware infiltration.

hot aisle In its simplest form, a type of data center design that involves lining up server racks in alternating rows with hot air exhausts facing one way and cold air intakes facing the other.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

hot site A near duplicate of the original site of an organization that can be up and running within minutes (maybe longer). Computers and phones are installed and ready to go, a simulated version of the server room stands ready, and the vast majority of the data is replicated to the site on a regular basis in the event that the original site is not accessible to users for whatever reason. Hot sites are used by companies that would face financial ruin in the case that a disaster makes their main site inaccessible for a few days or even a few hours. This is the only type of redundant site that can facilitate a full recovery.

hping A free TCP/IP packet generator, assembler, and analyzer that can be used to send large volumes of TCP traffic at a target while spoofing the source IP address, making it appear random or even originating from a specific user-defined source.

hybrid cloud A mixture of public and private clouds. Dedicated servers located within the organization and cloud servers from a third party are used together to form the collective network. In these hybrid scenarios, confidential data is usually kept in-house.

hybrid warfare A technique used to manipulate people's sentiment (often political or religious beliefs) with potentially false information or propaganda. Hybrid warfare is often done by leveraging social media sites such as Twitter, Facebook, Instagram, Reddit, and more.

Hypertext Transfer Protocol Secure (HTTPS) The secure version of HTTP, which is the primary protocol used to send data between a web browser and a website.

I

identification of critical systems The process of categorizing critical systems within an organization.

identity A distinguishing characteristic of a user, system, or application.

identity provider (IdP) The service provider that also manages the authentication and authorization process on behalf of the other systems in a federation.

identity theft The theft of your personal information, which is subsequently used to commit an act of fraud (credit card, taxes, medical) in your name.

IEEE 802.1X An IEEE standard that defines port-based network access control (PNAC). 802.1X is a data link layer authentication technology used to connect devices to a LAN or WLAN. It defines EAP.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

ifconfig A command used to configure kernel-resident network interfaces. It is used at boot time to set up interfaces as necessary. It is used to view TCP/IP configurations on a Linux or macOS system.

image backups Images of your entire operating system, including files, executable programs, and OS configurations. With an image backup, you can restore a single file, directory, or entire disk to the same or another device or to a virtual machine.

image steganography A technique used to hide any kind of file inside an image file.

impact The value and likelihood of a potential threat.

impact assessment A valuation to determine the potential monetary costs related to a threat.

impossible travel time/risky login A concept related to geolocation and geofencing where a potential compromise is identified based on the fact that it would be impossible for a user to be in two places at once. For instance, if you authenticated at 3 p.m. EST from New York and 10 minutes later tried to authenticate from Florida. This would not be possible and is an indication of account compromise.

incident response plan A set of instructions to help IT staff detect, respond to, and recover from network security incidents. These types of plans address issues like cybercrime, data loss, and service outages that threaten daily work.

incident response team A group of IT professionals in charge of preparing for and reacting to any type of organizational emergency.

incremental backups A series of backup data sets in which daily changes to the data are compared to the state of the data on the previous day. They all have to be applied to the original full backup copy to come up with an up-to-date full backup copy.

indicators of compromise Pieces of evidence or information that indicates a potential breach or compromise. Examples include command and control (C2) communications, IP addresses, domains, malware hashes, and other information.

industrial camouflage The use of blended images where the surrounding scenery and the camouflaged structure appear as one, with the goal to deceive passersby to believe the structure is something else entirely.

industrial control systems (ICS) Systems that monitor, control, sense, and warn engineers of all aspects of the processes in industrial plants.

information assurance The process of protecting data or information by preventing unauthorized modification.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

information lifecycle A four-step process of data collection that includes collecting data, storing data, determining how data is used, and disposing of the data.

Information Sharing and Analysis Centers (ISACs) Private-sector critical infrastructure organizations and government institutions that collaborate and share information between each other. ISACs exist for different industry sectors. Examples include automotive, aviation, communications, IT, natural gas, elections, electricity, financial services, health care, and many other ISACs.

infrastructure as a service (IaaS) A cloud service that offers computer networking, storage, load balancing, routing, and VM hosting. In this cloud solution, you rent infrastructure. You purchase virtual power to execute your software as needed. This is much like running a virtual server on your own equipment, except you are now running a virtual server on a virtual disk. This model is similar to a utility company model because you pay for what you use

infrastructure as code The process of managing and provisioning computer data centers through machine-readable definition files rather than physical hardware configuration or interactive configuration tools.

inherent risk The level of untreated risk in a process, system, or activity.

initialization vector (IV) attack An attack against a cryptographic implementation used to reverse encryption methods (such as RC4) and/or recover a pre-shared key (PSK). Examples of initialization vector attacks include attacks against legacy wireless implementations running WEP.

injection attacks A category of attacks where the threat actor "injects" malicious code or malicious traffic. Examples of injection attacks include SQL injection, command injection, cross-site scripting, and cross-site request forgery (CSRF or XSRF).

input validation A process that ensures the correct usage of data; it checks the data that is input by users into web forms and other similar web elements. If data is not validated correctly, it can lead to various security vulnerabilities, including sensitive data exposure and the possibility of data corruption.

instance awareness A concept used by cloud access security broker solutions to enforce policies on specific parts of an application or instance.

integer overflows Errors that occur when arithmetic operations in a program attempt to create a numeric value that is too big for the available memory space. They create a wrap and can cause resets and undefined behavior in programming languages such as C and C++. The security ramification is that the integer overflow can violate the program's default behavior and possibly lead to a buffer overflow.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

intellectual property (IP) theft The act of stealing intellectual property such as design documents, source code, and other sensitive information from an organization or individual.

intelligence fusion A technique to automatically extract threat intelligence data from heterogenous sources to analyze such data.

intent The intention or purpose of the attacker.

intermediate certificate authority In a certificate chain, the body that signs the end-entity certificate. It then handshakes with the root certificate, which represents the root certificate authority. It signs the intermediate certificate, which in and of itself is self-signing; this means that it not only creates the certificate but also signs it with its own private key.

internal actor Often referred to as an "insider"; an attacker who is already on the "inside" of the organization (it could be an employee or a contractor).

internal risk Risk to an organization from inside the environment—for instance, users or compromised legacy systems.

International Organization for Standardization (ISO) The standards body that was originally established in 2005 and later updated in 2013 to address the topics of organizational context, involvement of leadership, planning and objectives, support including resources and communication, operational aspects, evaluation of performance, and continuous improvement.

Internet Message Access Protocol (IMAP) An Internet protocol that allows you to access your email wherever you are, from any device.

Internet of Things (IoT) A system that is connected to a network or the Internet which is embedded with sensors, software, and other technologies. Examples of IoT systems include industrial control systems (ICS), sensors, thermostats, wearable systems, security cameras, and other systems exchanging data with other systems over the Internet.

Internet Protocol (IP) schema A standardized addressing/naming schema that helps you avoid conflicts within your on-premises network or cloud deployments and be able to correlate data among disparate systems.

Internet Protocol Flow of Information Export (IPFIX) An industry protocol created to collect and analyze network traffic flow information (metadata of the connections established between systems over a network).

intrusive The term used to categorize certain IP traffic or data that can potentially cause negative effects and even crash systems and applications.

IP scanners Tools that scan for IP addresses and various other information on the devices on your network. They are available for Windows and Linux.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

ipconfig A Windows command-line tool that displays all current TCP/IP network configuration values and refreshes Dynamic Host Configuration Protocol and Domain Name System settings.

IPFIX A standard used to collect and analyze network flow information from infrastructure devices such as network switches and routers.

IPsec A principle that defines the architecture for security services for IP network traffic. Also known as Internet Protocol Security or IP Security protocol.

isolation A policy by which different hardware and software technologies are designed to protect each process from other processes on the operating system.

J–K

jamming The process of sending unsolicited wireless signals to cause interference or a denial of service condition.

job rotation The process of rotating a particular task among a group of employees.

journalctl A Linux command-line tool used for viewing logs that are collected by systemd.

Kerberos An authentication protocol designed at MIT that enables computers to prove their identity to each other in a secure manner. It is used most often in a client/server environment; the client and server both verify each other's identity. This is known as two-way authentication or mutual authentication.

key escrow A process implemented to secure a copy of the user's private key (not the public key) in case it is lost.

key exchange Also referred to as key establishment; a method in cryptography by which cryptographic keys are exchanged between two parties, allowing use of a cryptographic algorithm.

key length The number of bits in an encryption algorithm's key.

key management The process of storing and exchanging cryptographic keys between systems.

key recovery agent A tool that you can set up to recover lost or corrupted keys when installing a certificate authority to a Windows Server.

key stretching A technique used to make a possibly weak key, typically a password or passphrase, more secure against a brute-force attack by increasing the resources (time and possibly space) needed to test each possible key.

keylogger Spyware that records your keystrokes.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

knowledge-based authentication (KBA) Authentication of an individual based on knowledge of information associated with his or her claimed identity in public databases.

known environment Penetration test environment where the tester starts out with a significant amount of information about the organization and its infrastructure.

L

lack of vendor support The situation that arises when a vendor no longer supports software or hardware updates, configuration assistance, and security patches of its products.

last known-good configuration The configuration of a system made just before a problem started.

lateral movement A post-exploitation technique, the main goal of which is to move from one device to another to avoid detection, steal sensitive data, and maintain access to these devices to exfiltrate the sensitive data. Lateral movement is also referred to as pivoting.

Layer 2 Tunneling Protocol (L2TP) A tunneling protocol used to connect virtual private networks. It does not include confidentiality or encryption on its own. It uses port 1701 and can be more secure than PPTP if used in conjunction with IPsec.

LDAP injection An attack similar to SQL injection; it uses a web form input box to gain access or exploits weak LDAP lookup configurations.

least privilege An approach by which subjects are given only the necessary privileges needed to do their intended job.

legacy platforms Technology solutions and platforms that are end-of-sale and end-of-support by a vendor or considered obsolete.

legacy systems Out-of-date systems that often can increase the risk to an environment based on its outdated software containing vulnerabilities.

legal hold A court order issued to maintain different types of data as evidence.

lightweight cryptography An encryption method that features a small footprint and/or low computational complexity. It is aimed at expanding the applications of cryptography to constrained devices such as the ever-expanding IoT market.

Lightweight Directory Access Protocol over SSL (LDAPS) A protocol used to read and write information to Active Directory. By default, LDAP traffic is transmitted unsecured, but you can enable LDAPS by using certificates.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

live boot media A USB flash drive or external hard disk drive containing a full operating system that can be booted to.

log aggregation The act of collecting logs from multiple systems in a network.

log collector Software that is able to receive logs from multiple sources (data input) and in some cases offers storage capabilities and log analysis functionality.

log files Special files that record significant events on your computer, such as when a user logs on to the computer or when a program encounters an error. Whenever these types of events occur, Windows records the event in an event log.

logger A shell command interface to the system log module.

logic bomb Code that has, in some way, been inserted into software; it is meant to initiate some type of malicious function when specific criteria are met.

logistics The detailed organization, implementation, and management of the flow of things between points.

M

MAC cloning attack An attack in which the attacker sniffs the network for valid MAC addresses and then uses those MAC addresses to perform other actions.

MAC flooding attack An attack that sends numerous unknown MAC addresses to a network switch to cause a DoS condition.

machine/computer certificate Typically, a type of certificate used when authenticating a machine or computer or user connecting to a network via a VPN or 802.1x. With a User certificate, the Subject Alternative Name field in the certificate contains the fully qualified domain name (FQDN), which is used in the authentication process.

macros A set of rules or patterns that specify how certain input could trigger a command, a series of commands, or any other operation in a system.

malicious flash drives Removable media (typically USB flash drives) used to transfer malware to a computer.

malicious USB cables Universal serial bus cables that inject keystrokes onto a victim's system upon being plugged into a USB-capable device.

malware Malicious software used by attackers to perform negative actions on a system or application.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

managed power distribution unit (PDU) Essentially, a power strip with multiple outputs designed to distribute electric power, especially designed to deliver power to racks of computers and networking equipment located within a data center.

managed security service providers (MSSPs) Companies that provide services to manage your security devices and that can also help monitor and respond to security incidents.

managed service provider (MSP) A cloud service model that can deliver network, application, system, and management services using a pay-as-you-go model.

managerial controls A category of controls that includes the techniques and concerns addressed by an organization's management (managers and executives). Generally, these controls focus on decisions and the management of risk. They also concentrate on procedures, policies, legal and regulatory, the software development lifecycle (SDLC), the computer security lifecycle, information assurance, and vulnerability management/scanning. In short, these controls focus on how the security of your data and systems is managed.

mandatory access control (MAC) An access control policy determined by a computer system, not by a user or owner, as it is in DAC. Permissions are predefined in the MAC model.

mandatory vacations *X* number of consecutive days' vacation over the course of a year as part of an employee's annual leave with an organization.

manual code review The process of reading source code line by line in an attempt to identify potential vulnerabilities.

mean time between failures (MTBF) The average number of failures per million hours of operation for a product in question.

mean time to failure (MTTF) The amount of time that an asset, system, or application takes before it fails.

mean time to repair (MTTR) The time needed to repair a failed device.

measured boot A feature introduced in Windows 8; it was created to help better protect your machine from rootkits and other malware.

memdump A memory dump; the process of taking all information content in RAM and writing it to a storage drive.

memorandum of understanding (MOU) A document that outlines the terms and details of an agreement between parties, including each party's requirements and responsibilities.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

memory leak A type of resource leak caused when a program does not release memory properly. This condition can cause an application to be vulnerable to exploitation by an attacker.

memory management The process of properly allocating and protecting memory in a system.

metadata Data created from every activity you perform, whether it's on your personal computer or online, every email, web search, social and public application. Metadata is defined as "data that provides information about other data."

MicroSD hardware security modules (HSMs) Physical devices that act as secure cryptoprocessors during secure login/authentication processes, during digital signings of data, and for payment security systems. HSMs can be found in various adapter card forms such as MicroSD HSM, as well as devices that plug into a computer via USB, and as network-attached devices.

microservices Applications that can be deployed as a collection of services that are highly maintainable and testable and independently deployable. They are organized around business capabilities and enable the rapid, frequent, and reliable delivery of large, complex applications. Most microservices developed and deployed by organizations are based on containers.

Mimikatz A tool used by many penetration testers, attackers, and even malware that can be useful for retrieving password hashes from memory; it is a useful post-exploitation tool.

MITRE ATT&CK A globally-accessible knowledge base of adversary tactics, techniques, and procedures (TTPs) based on real-world observations of cybersecurity threats.

mobile application management (MAM) A mobile device framework that can be used to address the types of applications that can be installed and which resources the applications can use.

mobile device management (MDM) Centralized software solutions that can control, configure, update, and secure remote mobile devices such as Android, iOS, BlackBerry, and so on, all from one administrative console. For networks with a lot of users, consider third-party offerings from companies that make use of MDM platforms.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

motion recognition The capability of newer cameras and software to recognize that something moved or changed position relative to its surroundings, or the surroundings to the object.

motivation The reason that drives a threat actor to perform a cyber attack against a system, individual, or organization.

multifactor authentication An authentication method that requires the user to provide two or more verification factors to gain access to a resource.

multifunction printer (MFP) An office machine that incorporates the functionality of multiple devices in one, so as to have a smaller footprint specialized in document management, distribution, and production in a large-office setting.

multiparty Involving more than one entity (individual or organization).

multipath Multiple physical routes established for storage. For example, multipath I/O defines more than one physical path between the CPU in a computer and its mass-storage devices through the buses, controllers, and bridge devices connecting them.

N

NarrowBand-Internet of Things (NB-IoT) An Internet of Things implementation based on low-power wide-area (LPWA) technology developed to enable a wide range of new IoT devices and services.

NAS (network-attached storage) backups Backups that are copied over the network to a storage array.

National Institute of Standards and Technology(NIST) An organization that is part of the United States Department of Commerce that is tasked to create different standards and technologies to help protect the United States critical infrastructure. NIST also works with private organizations and academia to develop such standards and guidance.

near-field communication (NFC) A communication method used by Bluetooth-equipped devices that enables two mobile devices (or a mobile device and a stationary computer) to be automatically paired and transmit data.

Nessus A remote security scanning tool that scans a computer and notifies the practitioner if it discovers any vulnerabilities that malicious hackers could use to gain access to any computer connected to a network.

netcat A back-end tool that allows for port scanning and port listening. In addition, you can actually transfer files directly through **netcat** or use it as a backdoor into other networked systems.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

NetFlow A session flow protocol that collects and analyzes network traffic data that can be used to help you understand which applications, users, and protocols might be consuming the most network bandwidth or if a DoS activity is taking place and who the actors are.

netstat A Windows and Linux command-line tool that generates a display showing network status and protocol statistics. It is used to view the current TCP/IP connections on a system.

network and port scanners Devices used to determine what TCP or UDP ports are open on the target system.

network DDoS attacks Attacks that target network infrastructure resources (for example, bandwidth, CPU, and memory utilization of the underlying network infrastructure).

network interface card (NIC) teaming The process of combining two or more network interfaces to increase network capacity.

network reconnaissance The testing of potential vulnerabilities in a computer network.

network redundancy The process of adding additional instances of network devices and connections to help ensure network availability and decrease the risk of failure.

network-based intrusion detection system (NIDS) A type of IDS that attempts to detect malicious network activities—for example, port scans and DoS attacks—by constantly monitoring network traffic.

network-based intrusion prevention system (NIPS) A type of IPS designed to inspect traffic and, based on its configuration or security policy, remove, detain, or redirect malicious traffic.

next-generation firewall (NGFW) A part of the third generation of firewall technology, combining a traditional firewall with other network device-filtering functions, such as an application firewall using in-line deep packet inspection or an intrusion prevention system.

nmap A popular vulnerability scanner.

nonce A random number issued by an authentication protocol that can be used only one time.

nonintrusive When IP traffic does not cause any negative effects in an application and underlying system.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

nonrepudiation The assurance that someone cannot deny the validity of something; where a statement's author cannot dispute its authorship.

normalization The capability to avoid or reduce data redundancies and anomalies—a core concept within relational databases.

nslookup A simple but practical command-line tool, that is principally used to find the IP address that corresponds to a host or the domain name that corresponds to an IP address.

null pointer dereference A situation that occurs when a program dereferences a pointer that it expects to be valid, but is null, which can cause the application to exit or the system to crash.

NXLog A tool that is used for centralized logging across various platforms and supports a myriad of different log types and formats.

O

OAuth A framework that provides authorization to a third-party entity (for example, a smartphone application) to access resources hosted on a resource server. In a classic client/server authorization framework, the third-party entity would receive the credentials from the resource owner (user) and then access the resource on the resource server.

obfuscation The complicating of source code to make it more difficult for people to understand. See also *camouflage*.

object detection A security camera functionality that enables cameras to evaluate and detect any new objects in a specific area, sometimes also employing facial recognition.

offboarding Removing an employee from a federated identity management system, typically when he or she leaves an organization.

offline password cracker A dedicated system or application used to crack hashed or encrypted passwords offline.

off-premises The act of storing and accessing data, creating applications, and interacting with thousands of other services over the Internet.

offsite storage backups The process of ensuring redundancy and protection by securing your data at a remote server away from your premises.

onboarding Adding a new employee to an organization and to its identity and access management system. This process is associated with user training, federated identity management, and role-based access control (RBAC).

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Online Certificate Status Protocol (OCSP) An alternative to certificate revocation lists (CRLs) that contains less information than a CRL does, and the client side of the communication is less complex. However, OCSP does not require encryption, making it less secure than CRLs.

online password cracker An application used to crack passwords while interacting with the targeted system.

online vs. offline backups An online backup places your files onto spare backup data servers. It is called an online backup because it uses the Internet to transfer files. The offline backup is a cold backup. A cold backup is performed while the database is offline and unavailable to its users.

on-path attack Previously known as man-in-the-middle (MITM) or man-in-the-browser (MITB) attack, this type of attack intercepts all data between a client and server, sometimes using a Trojan to infect a vulnerable web browser for later nefarious purposes.

on-premises Technical solutions deployed within an organization's infrastructure and not in the public cloud.

Opal An implementation profile for storage devices built to protect the confidentiality of stored user data against unauthorized access after it leaves the owner's control (involving a power cycle and subsequent deauthentication).

open-source intelligence (OSINT) Information that can be used for reconnaissance from public records, social media sites, DNS records, and other open sources. It applies to offensive security (ethical hacking/penetration testing) and defensive security. In offensive security, OSINT enables you to leverage public information from DNS records, social media sites, websites, search engines, and other sources for reconnaissance—in other words, to obtain information about a targeted individual or an organization. When it comes to threat intelligence, OSINT refers to public and free sources of threat intelligence.

Open Web Application Security Project (OWASP) A nonprofit organization that has chapters all over the world that focus on application and software security. It has numerous well-known and comprehensive projects designed to increase the awareness of secure coding and testing, as well as creating tools to help find and prevent security vulnerabilities.

OpenID A popular SSO protocol for federated systems. In the 2.0 version, the authentication and authorization process is similar to the one in SAML. OpenID also defines an IdP, called the OpenID provider (OP), and a relying party (RP), which is the entity that holds the resource the user wants to access. In OpenID, a user is free to select an OP of his or her choice, and the initial identity is provided in the form of a URL.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

OpenSSL A full-featured toolkit for the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols.

operational controls This category of controls includes the controls executed by people. They are designed to increase individual and group system security. They include user awareness and training, fault tolerance and disaster recovery plans, incident handling, computer support, baseline configuration development, and environmental security. The people who carry out the specific requirements of these controls must have technical expertise and understand how to implement what management desires of them.

operational technology (OT) The term used to describe physical items that can be programmed and connected to a network or the Internet. Typically, these devices are used to control electrical grids, pipelines, automobiles, manufacturing plant robots, and other critical infrastructure.

order of volatility The order in which digital evidence is collected from high volatility (where data is more vulnerable to loss) to low volatility.

organized crime A category of local, national, or international criminal groups of individuals engaged in illegal activity.

outsourced code development The development of code by a third-party organization.

P

P12/PFX A binary format based on PKCS#12 used to store a server certificate, intermediate certificates, and the private key in one encryptable file. It is typically used to import and export certificates and private keys.

packet capture The process of intercepting data packets that cross a specific point in a data network. It is sometimes used legitimately to troubleshoot a problem or to investigate a security issue.

pagefile A portion of a hard disk or similar media that is reserved as an extension of RAM.

partially known environment Penetration test environment where the penetration testers may be provided credentials but not full documentation of the network infrastructure.

pass the hash attack A type of attack in which, instead of trying to figure out what the user's password is, the attacker just uses a password hash collected from a compromised system and then uses the same hash to log in to another client or server system. This is done because password hashes cannot be reversed.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

passive reconnaissance An attacker method that can be carried out by researching information about the victim's public records, social media sites, and other personal information.

Password Authentication Protocol (PAP) A protocol that sends usernames and passwords in clear text. Obviously, this protocol is insecure and to be avoided. Therefore, PAP should not be used in networks today.

password cracker An application program that is used to identify an unknown or forgotten password to a computer or network resources.

password keys A technology typically deployed by corporations when implementing two-factor authentication. The primary use case is remote access to the organization's environment.

password reuse The practice of reusing the same password or part of it, which consequently increases the risk of password compromise.

password spraying A type of password brute-force attack where the attacker uses a single password against targeted user accounts before performing a second attempt to remain undetected.

password vault A central system or piece of software that stores and manages various sets of credentials in a secure management system. This password vault helps solve the issue of credential storage. The vault then has its own set of credentials and possibly another authentication factor that is used to access the vault when the credentials are needed for accessing a system. A password vault is also often referred to as a password manager.

patch management The process of keeping up with fixes that address software bugs.

patches Software bug fixes.

pathping A utility that sends packets to each router on the way to a final destination over a period of time and computes results based on the packets that return from each hop.

Payment Card Industry Data Security Standard (PCI DSS) A standard enacted in 2006 as a joint effort by the credit card industry with the overall goal to enhance the security around payment card data processing. The requirement applies to any organization that processes payment card data and enforces penalties for noncompliance on such organizations.

PEM Privacy-enhanced Electronic Mail (PEM); a common format that uses base64-encoded ASCII files and can be identified with the .pem file extension, though the format might also use .crt (for example, Microsoft), .cer, or .key extensions.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

penetration testing The process of analyzing the security posture of a network's or system's infrastructure in an effort to identify and possibly exploit any security vulnerabilities found.

perfect forward secrecy A feature of specific key agreement protocols that gives assurances that session keys will not be compromised even if long-term secrets used in the session key exchange are compromised.

persistence The act of maintaining a foothold in a compromised system after the exploitation phase in order to perform additional tasks such as installing and/or modifying services to connect back to the compromised system. It is used in a load-balancing environment, when various mechanisms are used to maintain the preservation of data during transmission between the client and server.

personally identifiable information (PII) Information used to uniquely identify, contact, or locate a person. This type of information could be a name, birthday, Social Security number, biometric information, and so on.

person-made disaster Disaster caused by the influence of humans.

personnel credential policy An organization's system of ensuring employees are who they say they are. The most common type of authentication is the username/password combination. Usernames are usually based on a person's real name.

pharming A type of attack in which an attacker redirects one website's traffic to another website that is bogus and possibly malicious. The threat actor redirects a victim from a valid website or resource to a malicious one that could be made to appear as the valid site to the user. From there, an attempt is made to extract confidential information from the user or to install malware in the victim's system.

phishing An attempt at fraudulently obtaining private information, usually done electronically. Vishing is done by phone. Spear phishing targets specific individuals. Whaling targets senior executives.

phishing campaigns In the context of this book and the Security+ blueprint, a method of user awareness training where the InfoSec organization employs a fake phishing attempt to educate its user base.

phishing simulations A method of user awareness training where the InfoSec organization employs a fake phishing attempt to educate its user base.

phone call authentication A verification system in which a user receives an automated phone call that requires him or her to press a certain button or code.

physical controls A physical security system used to control access to organization. It can be considered the first line of defense, sort of like a firewall is the first line of defense for a network. Implementing physical access security methods should

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

be a top priority for an organization. Proper building entrance access and secure access to physical equipment are vital. And anyone coming and going should be logged and surveyed.

physical locks Traditional locks designed to deter entry by unauthorized users who don't have a key.

piggybacking A type of attack in which an unauthorized person tags along with an authorized person to gain entry to a restricted area. Also known as tailgating.

ping A TCP/IP command used to verify IP-level connectivity to another TCP/IP computer by sending Internet Control Message Protocol (ICMP) echo request messages. Corresponding echo reply messages are displayed, along with round-trip times. This command is used to test connectivity between two devices on a network with IPv4.

pinning A method of adding security to the certificate validation process. You can help detect and block many types of on-path attacks by adding an extra step beyond normal X.509 certificate validation.

pivoting A post-exploitation technique, the main goal of which is to move from one device to another to avoid detection, steal sensitive data, and maintain access to these devices to exfiltrate the sensitive data. Also referred to as lateral movement.

plaintext Unencrypted data (typically an unencrypted password).

platform as a service (PaaS) A cloud service that provides various software solutions to organizations, especially the ability to develop applications in a virtual environment without the cost or administration of a physical platform. This model provides everything except applications. Services provided by this model include all phases of the system development lifecycle (SDLC) and can use application programming interfaces (APIs), website portals, or gateway software. These solutions tend to be proprietary, which can cause problems if the customer moves away from the provider's platform.

playbook Linear style electronic checklists of required steps and actions needed to successfully respond to specific incident types or threats.

port security A security feature present in routers and switches that is used to provide access control by restricting the Media Access Control (MAC) addresses that can be connected to a given port.

port spanning/port mirroring Configuring one or more ports on a switch to forward all packets to another port. This procedure is often used when capturing packets.

port tap A hardware device utilized to provide access to all traffic on a network segment.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Post Office Protocol (POP) An application-layer Internet standard protocol used by email clients to retrieve email from a mail server. POP version 3 is the version in common use.

post-quantum cryptography The use of cryptographic algorithms (usually public-key algorithms) that are thought to be secure against an attack by a quantum computer.

potentially unwanted programs (PUPs) A form of spyware that typically includes grayware, adware, or jokes.

PowerShell A scripting framework used in Windows operating systems.

predictive analysis A machine learning solution to help discover security threats in your network.

prepending Adding a message in an email's subject line to identify emails that come from outside the organization. You can often configure your email servers or email cloud services to use this technique.

preservation of evidence (evidence preservation) The act of making sure that digital evidence is acquired, handled, and analyzed properly and without any contamination or modification.

preshared key (PSK) A complex passphrase used to enable connectivity between wireless clients and the WAP. PSK is automatically used when you select WPA2-Personal in the Security Mode section.

pretexting Pretexting is the act of impersonating or "spoofing" someone else's identity.

preventive controls Controls that are employed before an event and are designed to prevent an incident. Examples include biometric systems designed to keep unauthorized persons out, NIPSs to prevent malicious activity, and RAID 1 to prevent loss of data. These are also sometimes referred to as deterrent controls.

Privacy Act of 1974 A federal law that establishes a code of fair information practice and governs the collection, use, and dissemination of personally identifiable information about persons' records maintained by federal agencies.

privacy notice A document sent from an organization that collects data stating how it is conforming to data privacy principles.

private cloud A type of cloud system designed for a particular organization in mind. As security administrator, you have more control over the data and infrastructure. A limited number of people have access to the cloud, and they are usually located behind a firewall of some sort in order to gain access to the private cloud.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Resources might be provided by a third party, or could come from your organization's server room or data center.

private information sharing centers Teams or groups of people who share threat intelligence information only within an organization and not publicly.

privilege escalation The process of elevating the level of authority (privileges) of a compromised user or a compromised application.

privileged access management (PAM) A system used to centrally manage access to privileged accounts. It's primarily based on the concept of least privilege.

production The environment where your code fulfills its destiny (where users access the final code after all of the updates and testing). When you hear people talk about making their "code go live," this is the environment they are talking about.

promiscuous mode Running a network interface or system in monitoring mode only.

proprietary Information that is critical and specific to an organization's business. This proprietary information is sometimes considered intellectual property.

protected cable distribution system (PDS) A system designed to deter, detect, and/or make difficult physical access to the communication lines carrying data and/or voice communications.

Protected Extensible Authentication Protocol (PEAP) An authentication protocol used to encapsulate Extensible Authentication Protocol (EAP) packets in order to safeguard sensitive data.

protocol analyzers Tools that allow network engineers and security teams to capture network traffic and perform analysis on the captured data to identify potential malicious activity or problems with network traffic.

provenance The establishment of a chain of custody for information that can describe its generation and all subsequent modifications that have led to its current state.

pseudo-anonymization A method of obfuscating sensitive data while not anonymizing the nonsensitive data so that it can be used for other business purposes.

public A classification of information where unauthorized access to the information does not cause any significant damage.

public cloud Application and storage space offered to the general public over the Internet by a service provider.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

public information sharing centers Organizations that share threat intelligence information publicly.

public key infrastructure (PKI) An entire system of hardware and software, policies and procedures, and people. It is used to create, distribute, manage, store, and revoke digital certificates. If you have connected to a secure website in the past, you have utilized a PKI.

public ledger Information organized into a long chain of blocks. When a buyer and seller engage in a transaction, the blockchain verifies the authenticity of their accounts. This is done by using the public ledger and by checking if the funds are available to proceed with the transactions.

public notifications and disclosures The act of providing the public notifications and information about a data breach.

pulping A data destruction method where paper is first shredded and then reduced to pulp.

pulverizing Grinding or shredding media and paper multiple times beyond recognition.

purple team A team that integrates the defensive capabilities of a blue team with the adversarial techniques used by the red team. In most cases, the purple team is not a separate team, but a solid dynamic between the blue and red teams.

push notifications When software triggers and provides certain information such as alerts, authentication attempts, updates, or any other notifications to a device (mobile device, laptop, or desktop) without the user deliberately requesting it.

Python A popular programming language used to develop applications in many environments.

Q

qualitative risk assessment An assessment that assigns numeric values to the probability of a risk and the impact it can have on the system or network.

quality assurance (QA) The maintenance and verification of a desired level of quality of software, a product, or service.

quantitative risk assessment Assessment that measures risk by using exact monetary values. It attempts to give an expected yearly loss in dollars for any given risk. It also defines asset values to servers, routers, and other network equipment.

quantum cryptography The science of exploiting quantum mechanical properties to perform cryptographic tasks.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

quarantine The action taken by some antivirus/antimalware scanners on a file infected with malware or a virus. When a file is quarantined, the file is moved to a location on disk where it cannot be executed.

R

race condition A difficult exploit to perform that takes advantage of the small window of time between when a service is used and its corresponding security control is executed in an application or operating system, or when temporary files are created.

radio frequency identification (RFID) A technology used to encode digital data in tags or smart labels. This digital data is captured by a reader via radio waves. RFID data from a tag or label is captured by a device that stores the data in a database.

rainbow table In password cracking, a set of precalculated encrypted passwords located in a lookup table.

ransomware A type of malware that restricts access to a computer system and demands a ransom be paid to restore access.

Raspberry Pi A low-cost, credit-card-sized computer that plugs into a computer monitor, keyboard, and mouse.

real-time operating system (RTOS) An operating system intended to serve real-time applications that process data as it comes in, typically without buffer delays. Processing time requirements are measured in tenths of seconds or shorter increments of time.

reception The public-facing part of a business or organization, where a reception desk places a friendly face in front of a visitor, essentially a buffer zone away from strangers accessing anyone in the company. Receptionists vet the person attempting to gain access to the facility or a particular employee.

recovery point objective (RPO) The acceptable latency of data, or the maximum tolerable time that data can remain inaccessible after a disaster.

recovery time objective (RTO) The acceptable amount of time to restore a function—for example, the time required for a service to be restored after a disaster.

red team A term used to refer to individuals who are performing adversarial simulation and penetration testing. However, a "true red team engagement" goes beyond the traditional scope of penetrating testing.

Redundant Array of Inexpensive Disks (RAID) A technology that ensures that a disk failure does not lead to lost data.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

refactoring Restructuring driver code.

reference architecture A methodology and set of tools created by the Cloud Security Alliance to help cloud providers with guidance on developing secure interoperability best practices.

registration authority (RA) The entity used to verify requests for certificates. If the request is deemed valid, the RA informs the certificate authority (CA) to issue the certificate.

registry A database of information, settings, options, and other values for software and hardware installed on all versions of Windows operating systems.

remote-access Trojan (RAT) A type of Trojan used to gain back-end access to a server, taking control of it, often for malicious purposes.

remote-access VPN A VPN type that enables organizations to establish VPN tunnels between two or more network infrastructure devices in different sites so that they can communicate over a shared medium such as the Internet. Many organizations use IPsec, GRE, and MPLS VPNs as site-to-site VPN protocols.

Remote Authentication Dial-In User Service (RADIUS) A remote authentication protocol that provides centralized administration of dial-up, VPN, and wireless authentication and can be used with EAP and 802.1X.

Remote Authentication Dial-in User Service (RADIUS) Federation An organization that has multiple RADIUS servers, possibly on different networks, that need to communicate with each other in a safe way. This is accomplished by creating trust relationships and developing a core to manage those relationships as well as the routing of authentication requests.

remote code execution (RCE) An attack that allows an attacker to obtain control of a target computer through some sort of vulnerability, thus gaining the power to execute commands on that remote computer at will.

replay attack A network attack in which a valid data transmission is maliciously or fraudulently repeated or delayed.

replication The process of copying data from one system to another so that a specific data set is in one or more locations.

reputation damage Damage to the reputation of an organization due to a security incident.

request for comments (RFC) An Internet Engineering Task Force (IETF) standard specification and documentation.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

residual risk The risk left over after a detailed security plan and disaster recovery plan have been implemented. After risk transference, risk avoidance, and risk reduction techniques have been implemented, an organization is left with a certain amount of residual risk.

resource exhaustion An attack against availability that is designed to bring the network, or access to a particular TCP/IP host/server, to its knees by flooding it with useless traffic.

resource policies Policies that govern the process of creating, assigning, and managing rules over the cloud resources that systems (virtual machines, containers, and so on) or applications use.

response and recovery controls Plans that you must have in place in the unfortunate event of a data breach, including the assembly of a team of experts within your organization, as well as legal counsel.

reverse proxy A proxy server that resides in front of the application servers (normally web servers) and functions as an entry point for Internet users who want to access the corporate internal web application resources.

revert to known state A process in which the desired outcome is to take a system back to a prior moment in time, or state of existence.

review logs System logs that are analyzed to find security problems.

right-to-audit In a network/services/data storage hosting agreement, a provision where the buyer has the right to examine the books and records of the seller to assure that the seller has complied with the contract, and that no employee of the buyer received any funds either directly or indirectly. This is used to ensure there is no fraud.

risk acceptance The willingness of an organization to accept a certain amount of risk. Sometimes, vulnerabilities that would otherwise be mitigated by the implementation of expensive solutions are instead dealt with when and if they are exploited. Also known as risk retention.

risk appetite The types and amount of risk, on a broad level, an organization is willing to accept in its pursuit of value.

risk assessment The attempt to determine the number of threats or hazards that could possibly occur in a given amount of time to your computers and networks.

risk avoidance The decision not to carry out a proposed plan because the risk factor is too great.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

risk awareness The capability of an organization to identify risks before they become a threat. The overall preparedness of an organization to mitigate risk.

risk control assessment The act of identifying and analyzing the potential risks in an organization. Risk control assessments can be done by internal stakeholders or by a third-party organization.

risk control self-assessment The act of performing an internal risk assessment without engaging a third-party auditor or organization.

risk management The identification, assessment, and prioritization of risks, and the mitigating and monitoring of those risks.

risk matrix/heat map A technique to visually represent the results of a risk assessment or to highlight the systems or data with a higher level of risk.

risk mitigation The process of reducing or eliminating a risk altogether.

risk register A tool created during risk assessment to track issues and address problems as they occur. Also known as a risk log.

risk transference The act of shifting, sharing, or "transferring" risk to a third party. An example of risk transference (also known as risk sharing) would be an organization that purchases cybersecurity insurance for a group of servers in a data center.

robot sentries Mechanized guards that provide a 24/7/365 watchman, continuously monitoring and alerting on differentials. Robot sentries report anything out of the normal to the manned security desk/office.

rogue APs Unauthorized wireless access points/routers that allow access to secure networks.

role-based access control (RBAC) An access model that, like mandatory access control (MAC), is controlled by the system, and, unlike discretionary access control (DAC), not by the owner of a resource. However, RBAC is different from MAC in the way that permissions are configured. RBAC works with sets of permissions, instead of individual permissions that are label-based.

role-based training A type of training where instructors and trainees act out the roles they might play, such as network administrator, security analyst, and so on.

root accounts The highest level of privilege available on a UNIX/Linux-based system.

root certificate authority The primary certificate authority server that generates digital certificates to users, devices, applications, or other subordinate certificate authorities.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

rootkit A type of software designed to gain administrator-level control over a computer system without being detected.

route A command used to view and manipulate the IP routing table in Linux and Microsoft Windows operating systems and also in other systems like routers and switches.

rsyslog A utility that builds on syslog capabilities by adding support for advanced filtering, configuration and output. Rsyslog and syslog-ng are similar.

rule-based access control An access model that defines whether access should be granted or denied to objects by comparing the object label and the subject label. Also known as label-based access control.

rules of engagement Documentation that specifies the conditions under which a security penetration testing engagement will be conducted.

runbook A series of conditional steps to perform actions, such as data enrichment, threat containment, and sending notifications automatically as part of the incident response or security operations process.

runtime error A program error that occurs while the program is running.

S

salt A unique, random string of characters known only to the site added to each password before it is hashed. Typically, this data is placed in front of each password.

sandboxing Using a security mechanism for separating running programs, usually in an effort to mitigate system failures and/or software vulnerabilities from spreading.

Sarbanes-Oxley A federal law enacted in 2002; it governs the disclosure of financial and accounting information.

scalability The property of a system to handle a growing amount of work by adding resources to the system.

scanless A Linux command-line utility for exploitation websites that performs port scans on your behalf. This tool helps in the early stages of penetration testing to run a port scan on multiple systems and have it not come back from your IP address.

screened subnet One or more routers or firewalls that provide controlled access to company assets like web servers. Previously known as demilitarized zone (DMZ).

script kiddie An unsophisticated individual with little or no skills when it comes to technology who maliciously uses code written by others that is freely accessible on the Internet.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

scripting The process of automating the execution of computer tasks. Automation and scripting greatly increase an organization's capability to detect and respond to threats.

SEAndroid A modification of SELinux made by Google to enable the use of SELinux on the Android platform. You might be familiar with Security Enhanced Linux (SELinux). This is a security enhancement for Linux that is based on Mandatory Access Control (MAC).

secrets management The management of API keys, passwords, and certificates through the use of some kind of secrets management tool. It provides a mechanism for managing the access and auditing the secrets used in the cloud environment.

secure cookies HTTP cookies that have a secure attribute set, which limits the scope of the cookie to "secure" channels.

Secure File Transfer Protocol (SFTP) A file protocol for transferring large files over the web.

Secure Real-Time Transport Protocol (SRTP) The secure version of RTP, intended to provide encryption, message authentication and integrity, and replay attack protection to the RTP data in both unicast and multicast applications.

Secure Shell (SSH) A network communication protocol that enables two computers or devices to communicate and share data.

Secure Web Gateway (SWG) A security solution that enables you to secure your remote workers' Internet access while not overloading the corporate Internet pipe. It is sometimes thought of as a cloud firewall. However, an SWG typically has many other protection mechanisms in place, including things like CASB. One example of an SWG is Cisco Umbrella.

Secure/Multipurpose Internet Mail Extensions (S/MIME) A technology that allows you to encrypt your emails using asymmetric cryptography.

security advisories Documents created by vendors, security researchers, and other organizations to disclose security vulnerabilities. Security advisories are also often called security bulletins by some vendors.

Security Assertion Markup Language (SAML) An open standard for exchanging authentication and authorization data between identity providers. SAML is used in many single sign-on (SSO) implementations.

Security Information and Event Management (SIEM) A specialized device or software used for security monitoring that collects, correlates, and helps security analysts analyze logs from multiple systems. The software works by collecting logs

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

and event data generated by an organization's application, security devices, and host systems and bringing it together into a single centralized platform, allowing companies to identify threats in real time.

security monitoring The process of monitoring network infrastructure, servers, endpoints, and applications to detect potential security problems and incidents.

Security Orchestration, Automation, and Response (SOAR) A set of solutions and integrations designed to allow organizations to collect security threat data and alerts from multiple sources. They provide automation and orchestration capabilities to help security operation center (SOC) analysts and security teams to detect and respond to security incidents.

security posture The combined security state of software, hardware, and data.

security posture assessment The act of verifying the combined state of software, hardware, and data security.

segmentation The separation of computing systems and/or data using different protocols and implementations.

self-encrypting drive (SED) A disk drive that uses an encryption key to secure the data stored on the disk.

self-signed certificates A document created and used on a system that not only creates the certificate but also signs it. It signs it with its own private key. This type of certificate does not offer the authentication usage that a publicly signed certificate would.

semi-authorized hacker An individual who does not have any affiliation with a company but risks breaking the law by attempting to hack a system and then notifying the administrator of the system that he or she was successful in doing so—just to let administrator know.

sensitive A classification of information where unauthorized access to the information could cause some damage to an organization. Examples of information or assets that could receive this type of classification are internal team email, financial information, and so on.

sensors Network devices that inspect network traffic.

sentiment analysis The use of tools and solutions to monitor customer sentiment and brand reputation.

separation of duties The division of tasks when more than one person is required to complete the tasks.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

serverless architecture The use of cloud platforms to host and/or to develop applications without the need to maintain the underlying infrastructure, the hypervisor, VMs, and operating systems. A good example of a serverless architecture is Amazon's AWS Lambda (https://aws.amazon.com/lambda/).

server-side execution The manipulation of an application, operating system, or hardware in a server typically by leveraging input validation vulnerabilities.

server-side request forgery (SSRF) A type of web application vulnerability where a web server or application does not sanitize input. This vulnerability enables an attacker to manipulate the application to make HTTP requests to a malicious site.

server-side validation The ability to properly handle application and user input to prevent a security vulnerability and server-side execution.

service accounts Accounts typically used on a server to provide a separate set of credentials and permissions to an application or service that is running. For instance, a server that is running Apache web server might have an "apache_user" account.

service-level agreement (SLA) Part of a service contract that formally and clearly defines exactly what a vendor is responsible for and what the organization is responsible for.

services integration A series of tools and technologies that are used to connect different systems, applications, code repositories, and physical or virtual network infrastructure to allow real-time exchange of data and processes.

session hijacking Exploiting a computer session in an attempt to gain unauthorized access to data, services, or other resources on a computer.

session replay A type of attack in which an attacker steals a user's valid session ID and reuses that ID to perform malicious transactions and activities with a web application.

sFlow This standard samples packets and provides analysis to identify unauthorized network activity and investigate DDoS attacks. The successor to NetFlow.

shadow IT An employee or a group of employees using IT systems, network devices, software, applications, and services without the approval of the corporate IT department.

shared and generic accounts User accounts that are shared between multiple users who share the same password. See also *credentials*.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

shell variables Variables that are contained exclusively within the shell in which they were set or defined. They are often used to keep track of ephemeral data, like the current working directory.

shimming Adding a small library that intercepts API calls.

Short Message Service (SMS) A mobile phone network facility that enables users to send and receive alphanumeric text messages of up to 160 characters on their cell phones or other handheld devices.

shoulder surfing Using direct observation to find out a target's password, PIN, or other such authentication information.

sideloading Loading third-party apps from a location outside the official application store for that device. Note that sideloading can occur in several ways: by direct Internet connection (usually disabled by default); by connecting to a second mobile device via USB OTG (USB On-The-Go) or Bluetooth; by copying apps directly from a microSD card; or via tethering to a PC or Mac.

Simple Network Management Protocol version 3 (SNMPv3) A network protocol that provides secure access to devices by authenticating and encrypting data packets over the network.

Simultaneous Authentication of Equals (SAE) A replacement for preshared key (PSK) that helps to protect against brute-force password attacks and offline dictionary attacks.

single loss expectancy (SLE) The loss of value in dollars based on a single incident.

single point of failure An element, object, or part of a system that, if it fails, causes the whole system to fail.

single sign-on (SSO) The ability to log in once but gain access to multiple systems without being asked to log in again.

site risk assessment The act of identifying and analyzing risk at a given location or site.

site survey Surveys typically performed using a variety of Wi-Fi analyzer tools to produce a heat map of all wireless activity in the area. They help determine the best placement for access points.

site-to-site VPN VPN type that enables users to work from remote locations such as their homes, hotels, and other premises as if they are directly connected to their corporate network.

skimming The method by which an attacker surreptitiously captures credit card information or information from other similar cards (gift cards, loyalty cards, identification cards, etc.) from a cardholder.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

smart card A device that falls into the category of "something a person has" and is known as a token. It's the size of a credit card and has an embedded chip that stores and transacts data for use in secure applications such as hotel guest room access, prepaid phone services, and more.

smart devices Devices ranging from A/C systems, garage doors, security systems, TVs, and so on that can be controlled for wired or wireless networks.

smart meters Electronic devices that record information such as consumption of electric energy, voltage levels, current, and power factor. Smart meters communicate the information of consumption behavior to the consumer, and electricity suppliers for system monitoring and customer billing.

smishing Using Short Message Service (SMS) to send malware or malicious links to mobile devices.

sn1per An automated scanner that you can use during a penetration test to perform vulnerability scanning.

snapshot backups Duplicate copies primarily used to restore a system, virtual machine, and disk or drive to an operational state and to serve as the system's restore point when the snapshot was taken.

social media Vulnerable sites like Twitter and Facebook that are used to influence the sentiment of a given user.

social media analysis A process that many organizations use to research a potential employee or volunteer. The procedure to do this includes investigating the person's online presence, including Facebook, Instagram, and LinkedIn.

software as a service (SaaS) A cloud service model that offers users access to applications that are provided by a third party over the Internet. It is designed to provide a complete packaged solution. The software is rented out to the user. The service is usually provided through some type of front end or web portal. While the end user is free to use the service from anywhere, the company pays a per-use fee. Examples of SaaS offerings are Office 365, Gmail, Webex, Zoom, Dropbox, and Google Drive.

software compliance/licensing Adherence to policies and standards with regard to the entitlement and legal use of software. Lack of software compliance and licensing can many times cause increased risk to an organization.

software-defined networking (SDN) A solution created to decouple control from the forwarding functions in networking equipment and to use software to centrally manage and "program" the hardware and virtual networking appliances to perform forwarding.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

software-defined visibility (SDV) A software-defined networking model designed and optimized to provide intelligent visibility of what's happening across an organization's network or in the cloud.

software development environments The platforms, infrastructure, and systems used to develop software and applications.

software diversity The use of different system variants that are developed or acquired simultaneously to address a wide range of business or customer requirements.

software integrity measurement The process of measuring your source code's quality when it is passed on to the quality assessment (QA).

spam Unwanted, unsolicited, "scam" email messages or social media posts.

spear phishing Using phishing to target specific groups of people or even specific individuals.

split tunneling Accessing a public network and a LAN at the same time using one or more network connections. A client system—for example, a mobile device—might do this.

spyware A type of malicious software either downloaded unwittingly from a website or installed along with some other third-party software.

SQL injection (SQLi) A vulnerability where an application does not sanitize user input and that enables an attacker to inject malicious SQL statements to potentially steal sensitive data, create, modify, or delete records in a database.

SSAE SOC 2 Type I/II Standards of Attestations Engagement (SSAE) System and Organization Controls (SOC); a standard whose main focus is related to the controls internal to an organization which are related to security compliance and operations. It is built on the concept of trust principles.

SSH *See* Secure Shell (SSH)

SSH keys The keys used to create a secure channel between two computers or network devices using SSH.

SSL Inspection (SSLI) The process of inspecting encrypted traffic. Typically, this is done by deploying a system acting as a proxy, decrypting, analyzing, and reencrypting the traffic. Modern implementations use the TLS protocol instead of SSL. Thus, this concept is now referred to as TLS inspection.

SSL stripping attack A man-in-the-middle (MITM) attack that can redirect a client to an insecure HTTP connection. The attacker still establishes a secure HTTPS connection between herself or himself and the server.

staging Testing your code or applications in a staging environment; however, this environment is as similar to the production environment as it can be.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

standard naming conventions Appropriate naming conventions for describing IT infrastructure, applications, and users to avoid conflicts and to be able to correlate data among disparate systems.

stapling Allowing the presenter of a certificate to bear the cost involved when providing Online Certificate Status Protocol (OCSP) responses. Previously known as TLS Certificate Status Request.

state actors A government using cyber-attack techniques for different purposes, such as stealing information, reconnaissance, espionage, and disruption.

static code analysis The analysis of computer software that is performed without actually executing programs.

steganography The practice of hiding a secret message inside of or even on top of something that is not secret, such as an image, video, or audio file.

storage area networks (SANs) The most common storage networking architecture used by enterprises for business-critical applications that need to deliver high throughput and low latency. By storing data in centralized shared storage, SANs enable organizations to apply consistent methodologies and tools for security, data protection, and disaster recovery.

stored procedure A combination of precompiled SQL statements, stored in the database, that execute some task. They combine machine learning (ML) with automation to respond to threats and maintain critical operations.

strategic intelligence The strategy of narrowing down an investigation to relevant information gathering initially, which is the use of all resources to make determinations that limit an investigation's scope to a manageable level.

stream cipher A way of combining plaintext digits with a pseudorandom cipher digit stream, or keystream.

structured exception handling (SEH) A mechanism used to handle software and hardware exceptions.

Structured Threat Information eXpression (STIX) An OASIS standard designed for sharing of cyber-attack information in a machine-readable format (JSON).

Subject Alternative Name (SAN) A field in which an organization can specify additional hostnames, domain names, IP addresses, and so on.

subscriber identity module (SIM) Widely known as a SIM card; an integrated circuit that is intended to securely store the international mobile subscriber identity

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

(IMSI) number and its related key, which are used to identify and authenticate subscribers on mobile telephony devices like mobile phones and computers.

supervisory control and data acquisition (SCADA) Systems capable of managing parts inventories for just-in-time manufacturing, regulate industrial automation and robots, and monitor process and quality control.

supply chain The process, people, and activities involved in creating and supplying a product or service to a consumer. In the context of this book and the Security+ blueprint, the specific concerns for a supply chain when addressing third-party risk management.

supply-chain attack A type of attack in which attackers target security weaknesses in the supply network and install malicious software or hardware implants to perform different nefarious activities.

swap file An area of the hard drive that can be used by the operating system to store data from RAM that has not been used in recent activity, typically moving data/apps required from disk to RAM and from RAM back to disk, hence the name "swapping."

symmetric encryption An encryption algorithm implementation that uses a single key that needs to be shared among the people who need to receive the message.

syslog A standard for message logging. It is available on most network devices (e.g., routers, switches, and firewalls), as well as printers, and UNIX/Linux-based systems.

syslog-ng A log management solution that helps improve the performance of a SIEM system.

system integration The process of incorporating different components into one system and making sure that they operate correctly.

system on a chip (SoC) An integrated circuit that integrates all or most components of a computer system or other electronic system to perform a complete function.

T

tabletop exercise A practice that begins with a security incident preparedness activity, taking participants through a step-by-step process of dealing with a simulated incident scenario and providing hands-on training for participants that can then highlight flaws in incident response planning.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

tag In digital forensics, a document that contains general information about an item and the incident under investigation and also records the date, time, control number, and name or initials of the investigator.

tail A Linux command that allows you to display the last 10 lines of any text file.

tailgating A situation in which an unauthorized individual follows an authorized individual to enter a restricted building or facility.

tape backups Magnetic tape or tape cartridges used as storage devices that can hold enormous amounts of data.

Tcpdump A data-network packet analyzer computer program that runs under a command-line interface.

Tcpreplay A suite of free open-source utilities for editing and replaying previously captured network traffic.

technical controls A category of controls that includes the logical controls executed by the computer system. Technical controls include authentication, access control, auditing, and cryptography.

temporal group A Common Vulnerability Scoring System (CVSS) group that assesses the vulnerability as it changes over time.

Terminal Access Controller Access-Control System Plus (TACACS+) A remote authentication protocol developed by Cisco. TACACS+ uses inbound port 49 like its forerunners; however, it uses TCP as the transport mechanism instead of UDP.

terms of agreement A document that protects personal information and the individual whose data is included.

testing A dedicated environment to help you make sure that your code does its job.

theft of intellectual property Theft of such information is a primary goal of internal and external threat actors.

thin clients Computer systems that run from resources stored on a central server or from the cloud instead of a local (on-premises) system. When you use a thin client, you connect remotely to a server-based computing environment where the applications, sensitive data, and memory are stored.

third-party libraries and software development kits (SDKs) Third-party libraries are applications, binaries, or pieces of code that are developed by a third-party organization or open-source software. SDKs are a collection of software development tools and utilities used as a framework to develop applications.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

threat feeds Threat intelligence sources from open sources and commercial vendors.

threat hunting Proactively and iteratively looking for threats in your organization.

threat intelligence The knowledge about an existing or emerging threat to assets, including networks and systems.

threat likelihood The probability that a threat might occur and be leveraged by an attacker to perform some action.

threat map A real-time map of security attacks happening at any given time. Threat maps have been used to also illustrate and visualize global cyberattacks. Also known as an attack map.

time-based one-time password (TOTP) A one-time password generated by taking its uniqueness from the current time.

time of check (TOC) or time of use (TOU) A race condition attack that takes advantage of the small window of time between when a service is used and its corresponding security control is executed in an application or OS, or when temporary files are created. It can be defined as anomalous behavior due to a dependence on timing of events.

token Hardware device or one implemented in software (soft-token) that essentially acts as a one-time password generator. A token can also be a cryptographic unique identifier of an application requesting access to a given service.

token key A cryptographic key of an application requesting access to a given service or resource.

tokenization A form of data obfuscation that is used in privacy enhancing technology. The approach of tokenization is to utilize a token to replace or obscure the data in a reversible manner. The process randomly generates a token value for plaintext data and stores the mapping in a database.

tracert/traceroute Computer network diagnostic commands for displaying possible routes and measuring transit delays of packets across a network, including the Internet. **tracert** is used on Windows systems; **traceroute** is used on Linux and macOS systems.

transit gateway A system that can be used to interconnect a virtual private cloud (VPC) and on-premises networks. Transit gateways can also be used to connect cloud-hosted applications to software-defined wide area network (SD-WAN) devices.

Transport Layer Security Inspection (TLSI) A security process that allows organizations to decrypt traffic, inspect the decrypted content for threats, and then reencrypt the traffic before it enters or leaves the network.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

transport mode A IPsec mode in which two hosts establish a VPN connection. Transport mode is allowed only between two hosts and not more. IPsec devices configured in transport mode only encrypt the payload and ESP trailer. The IP header of the original packet is not encrypted.

Trojan An application that appears to perform desired functions but is actually performing malicious functions behind the scenes.

true negative A system response that occurs when the intrusion detection device identifies an activity as acceptable behavior and the activity is actually acceptable.

true positive The successful identification of a security attack or a malicious event.

Trusted Automated eXchange of Indicator Information (TAXII) An OASIS standard that provides an open transport mechanism that standardizes the automated exchange of cyberthreat information. TAXII was originally developed by MITRE and is now maintained by OASIS.

trusted operating system (TOS) An operating system that has support for multilevel security and can meet government requirements. Examples of certified trusted operating systems include Windows 10, SELinux, FreeBSD (with the TrustedBSD extensions), and Red Hat Enterprise Server. To be considered a TOS, the manufacturer of the system must have strong policies concerning updates and patching.

Trusted Platform Module (TPM) (1) Per NIST, a tamper-resistant integrated circuit built into some computer motherboards that can perform cryptographic operations (including key generation) and protect small amounts of sensitive information, such as passwords and cryptographic keys. (2) an international standard for a secure crypto-processor, a dedicated microcontroller designed to secure hardware through integrated cryptographic keys.

tunnel mode A mode that enables hosts behind one secure gateway to communicate securely with hosts behind another gateway.

two-person integrity/control An escort and an observer that are assigned to a particular person doing work in a high-risk area, such as a data center, where a single person can cause massive amounts of damage in a few minutes.

typo squatting A technique used by adversaries to leverage user errors (typos) when entering a given URL in their web browser for a given website. An attacker can register the "misspelled" domain to impersonate the real website, host malware, or perform any other malicious technique to compromise the end user's system.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

U

unauthenticated mode An encryption protocol that provides no authentication of the encrypted data.

unauthorized hacker Malicious individual who attempts to break into computers and computer networks without authorization.

unified endpoint management (UEM) An approach that enables administrators to manage end-user mobile devices by controlling the deployment, operations, and monitoring of mobile devices used to access organization resources. It is used to enforce an organization's security policy on mobile devices. UEM also has the capabilities of a Mobile Application Management (MAM) system, Mobile Content Management (MCM), Mobile Threat Management (MTM), Containerization, and Identity and Access Management (IAM).

Unified Extensible Firmware Interface (UEFI) A specification for a software program that connects a computer's firmware to its operating system (OS).

unified threat management (UTM) A security product that evolved from the firewall and also includes IDS/IPS, antivirus, VPN, content filtering, DLP, and load balancing, among other technologies.

unknown environment A penetration test environment where the tester is typically provided only a very limited amount of information. For instance, the tester may be provided only the domain names and IP addresses that are in scope for a particular target. The idea of this type of limitation is to have the tester start out with the perspective that an external attacker might have.

URL redirection attack A type of attack in which an attacker redirects a user to a malicious site (URL) to install malware, to steal sensitive information, or to manipulate the user.

USB data blocker A device that stops an attacker from connecting a USB drive to a computer and stealing data.

user account An account assigned to an individual user for identity and authentication purposes.

user behavior analysis A type of security evaluation used to potentially find insiders who might be internal attackers.

user training User security awareness training.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

V

vendor management The process of managing third-party vendors. Third-party vendors include outsourcing companies, the software and hardware you purchase for your infrastructure, cloud service providers, and more. Risk management should include the evaluation of third-party vendors, their security capabilities, and even including security requirements in contract negotiations. In addition, if you are employing contractors from an external vendor, you should always audit the level of access that those contractors have to your resources and data.

vendors Companies that sell security hardware, software, or services.

version control The systems and process used to manage, track, and maintain software versions. In software development, the process of checking code in to a version control repository (for example, GitLab or GitHub), where each change is integrated and tested with the rest of the software system.

video steganography A technique to hide any kind of file in a cover video file.

virtual desktop infrastructure (VDI) A technology that enables users to connect to a virtual host (desktop) typically in a data center or in the cloud.

virtual machine (VM) A computer file, typically called an image, that behaves like an actual computer.

virtual private cloud (VPC) endpoint A virtual endpoint device that is able to access VPC services in a cloud computing environment. There are various types of VPC endpoint devices that can be utilized in an environment, based on the vendor.

virtual private network (VPN) A connection between two or more computers or devices not on the same private network.

virtualization A solution that allows you to create IT services using hardware resources that can be shared among different installations of operating systems and applications.

vishing Voice phishing; a type of attack in which attackers typically impersonate and spoof caller ID to obfuscate themselves in regular phone conversations to steal sensitive information such as Social Security numbers, credit card numbers, or other information used in identity theft schemes.

Visual Basic for Applications (VBA) An event-driven programming capability in Microsoft operating systems and applications.

VM escape attack An attack against a virtual technology platform where an attacker or malware compromises one VM and then attacks the hypervisor and potentially compromises other VMs.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

VM sprawl A situation that occurs when an organization can no longer effectively control and manage all the VMs on a network or in the cloud.

Voice over Internet Protocol (VoIP) A process for sending audio signals, primarily voice, over a data network, such as the Internet.

VPN concentrator Hardware or software used to terminate remote access or site-to-site VPN connections.

vulnerability databases A collection of vulnerabilities, often identified by Common Vulnerability and Exposure (CVE) identifiers.

W

war driving Collecting information about wireless networks by just driving or walking around.

war flying Flying drones or similar devices to obtain information about a wireless network or even collecting pictures of facilities in some cases.

warm site A work site that has computers, phones, and servers, but they might require some configuration before users can start working on them. The warm site has backups of data that might need to be restored; they are probably several days old. This solution is chosen the most often by organizations because it has a good amount of configuration yet remains less expensive than a hot site.

watering hole attack A type of attack in which an attacker profiles the websites that the intended victim accesses. The attacker then scans those websites for possible vulnerabilities. If the attacker locates a website that can be compromised, the website is then injected with JavaScript or other similar code injection that is designed to redirect the user when the user returns to that site (also known as a pivot attack). The user is then redirected to the malicious site.

wearables Devices such as smart watches, pedometers, glasses, 3D headsets, and smart clothing that are Internet-aware.

web application firewall (WAF) An application firewall used to protect servers (and their client sessions) from XSS and SQL injection, among other attacks, during HTTP sessions.

web application scanners Vulnerability scanners used to assess web applications and web services (such as APIs).

web of trust A decentralized trust model that addresses issues associated with the public authentication of public keys common to CA-based PKIs.

whaling An attack similar to phishing and spear phishing; however, with whaling, the attack is targeted at high-profile business executives and key individuals in a corporation.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

white team A term used to define a team that blends all previous colors together. White teams are individuals who are focused on governance, management, risk assessment, and regulatory compliance.

white-box testing A way of testing the internal workings of an application or system where the tester has knowledge of and detailed information about the system being tested.

Wi-Fi analyzers Hardware or software designed to give you a picture of the wireless signal and channel saturation in the area you are surveying.

Wi-Fi disassociation attack A type of attack in which the attacker deauthenticates a user from a wireless network. This attack is also referred to as a Wi-Fi deauthentication attack.

Wi-Fi Protected Access 2 (WPA2) The next evolution of Wi-Fi Protected access that was created to enhance and replace the original Wi-Fi Protected Access protocol or WPA. WPA2 includes support for Counter-mode/CVC-MAC protocol (CCMP) and AES-based encryption.

Wi-Fi Protected Access 3 (WPA3) A more robust authentication mechanism than WPA2. It also provides a higher level of encryption capabilities and enables a very robust authentication based on passwords by utilizing a technology called Simultaneous Authentication of Equals (SAE), which helps protect against brute-force password attacks and offline dictionary attacks.

Wi-Fi Protected Setup (WPS) A security standard originally created to give users easy connectivity to a wireless access point. It is a security vulnerability and should be disabled due to its vulnerability to brute-force attacks.

wildcard certificate A single certificate that allows connections to a main website and its subdomains.

WinHex A powerful application that functions as an advanced hex editor; a tool for data analysis, editing, and recovery; a data editing and wiping tool; and a forensics tool used for evidence gathering.

wireless access point (WAP) placement The process of determining the best placement for access points, typically in the center of a building to provide equal access to everyone on the perimeter of the organization's property, and there is the least chance of the signal bleeding over to other organizations.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Wireshark A widely used network protocol analyzer that enables packet captures on wireless, Bluetooth, USB, and LAN network-attached cards, placing them into what's known as promiscuous mode. Wireshark is used to troubleshoot network problems, examine security problems, verify network applications, debug protocol implementations, and learn network protocols present on networks.

worm Code that runs on a computer without the user's knowledge; a worm self-replicates, whereas a virus does not.

X–Z

X.509 A common PKI standard developed by the ITU-T that often incorporates the single sign-on (SSO) authentication method. Most certificates are based on the this standard.

XML injection An attack that can compromise the logic of Extensible Markup Language (XML) applications—for example, XML structures that contain the code for users. It can be used to create new users and possibly obtain administrative access.

zero-day vulnerability A type of vulnerability that is disclosed by any individual or exploited by an attacker before the creator of the software can create a patch to fix the underlying issue. Attacks leveraging zero-day vulnerabilities can cause damage even after the creator knows of the vulnerability because it may take time to release a patch to prevent the attacks and fix damage caused by them.

Zigbee A specification that uses the 2.4 GHz band and a self-healing true mesh network; it was created on the IEEE's 802.15.4 standard. Zigbee is intended as simpler and less expensive wireless personal area network (WPAN) alternative to Bluetooth and Wi-Fi.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

This page intentionally left blank

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Answers to the "Do I Know This Already?" Quizzes and Review Questions

Chapter 1

Do I Know This Already?

1. A. Spear phishing is one of the most common social engineering attacks where the attacker searches for public information about the victim to send a targeted email to steal information. Typo squatting (or typosquatting) is a technique used by adversaries that leverages human error when typing a URL in their web browser. Pharming is the term used to describe a threat actor redirecting a victim from a valid website or resource to a malicious one that could be made to appear as a valid site to the user. From there, an attempt is made to extract confidential information from the user or to install malware in the victim's system.
2. B. The Social Engineering Toolkit (SET) is an example of a tool that can be used specifically to perform social engineering attacks.
3. A. Vishing is social engineering attack in which the attacker calls the user over the phone and then persuades the user to reveal sensitive information or perform a given action. Smishing is a type of phishing campaign using SMS text messages instead of email.
4. D. An access control vestibule is a small space that can usually fit only one person, used to combat tailgating. Tunnel-gap or tunnel-traps are not correct social engineering terms. Piggyback is the act of following someone while opening a door to enter a building or a room.
5. A. Pretexting is the act of impersonating someone else.
6. A. Malvertising is a social engineering technique where an attacker incorporates malicious ads on trusted websites, which results in users' browsers being inadvertently redirected to sites hosting malware.
7. B. Spear phishing is phishing attempts that are constructed in a very specific way and directly targeted to specific individuals or companies.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

8. C. Whaling is a social engineering attack similar to phishing and spear phishing. However, in whaling attacks the attacker targets executives and key personnel of an organization (aka the "big fish").
9. B. Attackers use the social engineering scarcity technique to create a feeling of urgency in a decision-making context. It is possible to use specific language in an interaction to present a sense of urgency and manipulate the victim.
10. D. All of the available answers can be used as recommendations for user security awareness training and education.

Review Questions

1. Dumpster diving
2. Social engineering
3. Prepending
4. Lack of user awareness
5. A public building with shared office space
6. Typo squatting
7. Social engineering
8. Tailgating
9. Pharming
10. To deter shoulder surfing

Chapter 2

Do I Know This Already?

1. A. Ransomware is a type of malware that restricts access to a computer system and demands that a ransom be paid. It informs the user that in order to decrypt the files or unlock the computer to regain access to the files, a payment would have to be made to one of several banking services (typically crypto currencies like Bitcoin).
2. A. Trojans appear to perform desirable functions but are actually performing malicious functions behind the scenes.
3. B. Rootkit is a type of malware designed to gain administrator-level control over a computer system without being detected.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

4. C. Fileless malware works differently from traditional malware that puts malicious executables within the file system; instead, it works in a memory-based environment.
5. A. A group of compromised computers (bots), known as a botnet, is typically controlled by a command-and-control (C2) server/system.
6. B. A dictionary password attack pulls words from the dictionary or word lists to attempt to discover a user's password. A dictionary attack uses a predefined dictionary to look for a match between the encrypted password and the encrypted dictionary word.
7. A. In password spraying an attacker attempts to compromise a system using a large number of usernames with a few commonly used passwords.
8. A. Skimming is a type of attack in which an attacker captures credit card information or information from other similar cards (gift cards, loyalty cards, identification cards, and so on) from a cardholder surreptitiously. Attackers use a device called a skimmer that can be installed at strategic locations such as ATMs and gas pumps to collect card data.
9. D. Tainting, overfitting, and transfer attacks are types of adversarial techniques against machine learning (ML) implementations.
10. A. A supply-chain attack occurs when attackers target security weaknesses in the supply network and install malicious software or hardware implants to perform different nefarious activities.
11. D. Attackers can perform virtual machine (VM) escape, API, and DNS attacks to compromise cloud-hosted applications and services.
12. B. A downgrade attack is a type of cryptographic attack that forces the rollback of a strong algorithm in favor of an older, lower-quality algorithm or mode of operation.

Review Questions

1. Botnet
2. Bot
3. Pop-up windows with advertisements
4. You have been infected with a worm.
5. Ransomware
6. Logic bomb
7. Through email
8. Trojan

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Chapter 3

Do I Know This Already?

1. A. The two types of privilege escalation attacks are vertical and horizontal. A horizontal privilege escalation attack occurs when a user accesses functions or content reserved for other users. Vertical privilege escalation occurs when a lower-privileged user accesses functions reserved for higher-privileged users—for example, if a standard user can access functions of an administrator. This is also known as privilege elevation and is the most common description. To protect against this type of situation, you should update the network device firmware. In the case of an operating system, it should again be updated, and use of some type of access control system is also advisable—for example, User Account Control (UAC).

2. B. Stored, or persistent, XSS attacks occur when the malicious code or script is permanently stored on a vulnerable or malicious server, using a database. These attacks are typically carried out on websites hosting blog posts (comment forms), web forums, and other permanent storage methods. An example of a stored XSS attack is a user requesting the stored information from the vulnerable or malicious server, which causes the injection of the requested malicious script into the victim's browser. In this type of attack, the vulnerable server is usually a known or trusted site.

3. A. DLL injection occurs when code is run within the address space of another process by forcing it to load a dynamic link library (DLL). Ultimately, this type of attack can influence the behavior of a program that was not originally intended. This attack can be uncovered through penetration testing.

4. D. A null pointer dereference occurs when a program dereferences a pointer that it expects to be valid but is null, which can cause the application to exit or the system to crash. From a programmatical standpoint, the main way to prevent this situation is to use meticulous coding. Programmers can use special memory error analysis tools to enable error detection for a null pointer dereference.

5. D. Directory traversal, path traversal, and the ../ ("dot-dot-slash") attack are methods of accessing unauthorized parent (or worse, root) directories. They are often used on web servers that have PHP files and are Linux or UNIX-based but can also be perpetrated on Microsoft operating systems (in which case, it would be ..\ or the "dot-dot-backslash" attack).

6. B. Integer overflows occur when arithmetic operations attempt to create a numeric value that is too big for the available memory space. This creates a wrap and can cause resets and undefined behavior in programming languages

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

such as C and C++. The security ramification is that the integer overflow can violate the program's default behavior and possibly lead to a buffer overflow.

7. B. Race conditions are also known as time-of-check (TOC) or time-of-use (TOU) attacks.

8. C. Error handling or error exception handling code should be checked thoroughly so that a malicious user can't find out any additional information about the system. These error handling methods are sometimes referred to technically as pseudocodes.

9. B. A fuzzer is a program that can send crafted messages to a vulnerable application or system to find input validation vulnerabilities.

10. D. A replay attack is a network attack in which a valid data transmission is maliciously or fraudulently repeated or delayed. It differs from session hijacking in that the original session is simply intercepted and analyzed for later use. In a replay attack, an attacker might use a packet sniffer to intercept data and retransmit it later. In this way, the attacker can impersonate the entity that originally sent the data.

11. A. Session replay attacks occur when an attacker steals a user's valid session ID and reuses that ID to perform malicious transactions and activities with a web application.

12. D. Cross-site request forgery (XSRF) attacks leverage the trust that the application has in the targeted user. For example, the attacker could inherit the privileges of the user to perform an undesired action, such as stealing sensitive information, creating users, or downloading malware.

13. D. All of the available answers are best practices to help protect application programming interfaces (APIs).

14. C. Resource exhaustion is an attack against availability that is designed to bring the network, or access to a particular TCP/IP host/server, to its knees by flooding it with useless traffic. Resource exhaustion attacks are a form of denial-of-service (DoS) attacks. They can also leverage software vulnerabilities such as memory leaks and file descriptor leaks.

15. D. A memory leak is a type of resource leak caused when a program does not release memory properly. The lack of freed-up memory can reduce the performance of a computer, especially in systems with shared memory or limited memory. A kernel-level leak can lead to serious system stability issues. The memory leak might happen on its own due to poor programming, or it could be that code resides in the application that is vulnerable and is later exploited by an attacker who sends specific packets to the system over the network.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

16. A. An attacker can launch an SSL strip attack in different ways. One of the most common ways is to create a wireless hotspot and lure the victims to connect to it.
17. D. Attackers can potentially modify drivers through the use of driver shimming (the adding of a small library that intercepts API calls) and driver refactoring (the restructuring of driver code).
18. D. Pass the hash attacks leverage deficiencies in Windows NTLM implementations.

Review Questions

1. Mimikatz
2. Driver refactoring
3. SSL stripping attack
4. Swagger
5. Clickjacking
6. Race condition
7. Address space layout randomization (ASLR)
8. XML External Entity (XXE)

Chapter 4

Do I Know This Already?

1. A. An initialization vector (IV) attack is a type of related-key attack, which occurs when an attacker observes the operation of a cipher using several different keys and finds a mathematical relationship between those keys, allowing the attacker to ultimately decipher data.
2. B. ARP cache poisoning (also known as ARP spoofing) is an example of an attack that leads to a man-in-the-middle scenario. An ARP spoofing attack can target hosts, switches, and routers connected to a Layer 2 network by poisoning the ARP caches of systems connected to the subnet and by intercepting traffic intended for other hosts on the subnet.
3. A. With an on-path attack, the attacker must first infect the victim's computer with a Trojan. The attacker usually gets the malware onto the victim's computer through some form of trickery or deceit.
4. D. Attackers launch MAC flooding attacks by sending numerous unknown MAC addresses to a network switch to cause a DoS condition. In addition,

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

when the Layer 2 forwarding table limit is exceeded, packets are flooded to all ports in a virtual LAN (VLAN). This, in turn, enables the attacker to sniff network connections over a switched network while disrupting network performance.

5. D. Domain hijacking is a type of hijacking attack in which the attacker changes the registration of a domain name without the permission of the original owner/registrant. One of the most common methods to perform a domain hijacking is using social engineering.
6. B. One specific type of DDoS attack is the DNS amplification attack. Amplification attacks generate a high volume of packets ultimately intended to flood a target website.
7. B. Visual Basic for Applications (VBA) is an event-driven programming capability in Microsoft operating systems and applications. Attackers have used VBA to create malicious macros in applications such as Excel or Word.

Review Questions

1. PowerShell
2. DMARC
3. Cross-site scripting (XSS)
4. DNS poisoning
5. MAC cloning or spoofing

Chapter 5

Do I Know This Already?

1. A. The name *hacktivist* is often applied to different kinds of activities—from hacking for social change, to hacking to promote political agendas, to full-blown cyberterrorism. Due to the ambiguity of the term, a hacktivist could be inside a company or attack from the outside and will have varying amounts of resources and funding. However, a hacktivist is usually far more competent than a script kiddie.
2. B. Cybercriminals might work on their own, or they might be part of criminal syndicates and organized crime—a centralized enterprise run by people motivated mainly by money.
3. D. The level of sophistication/capability, resources, and funding are all attributes of threat actors that put them into different categories (that is, state-sponsored actors, script kiddies, hacktivists, criminals).

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

4. C. The Diamond Model is designed to represent a cybersecurity incident and is made up of four parts. Active intrusions start with an adversary who targets a victim. The adversary will use various capabilities along some form of infrastructure to launch an attack against the victim. Capabilities can be various forms of tools, techniques, and procedures, while the infrastructure is what connects the adversary and victim. The lines connecting each part of the model depict a mapping of how one point reached another. This mapping helps you understand the motives, intent, sophistication, capabilities, and resources that a threat actor may have.
5. A. One of the most effective attacks for mass compromise is to attack the supply chain of a vendor to tamper with hardware and/or software. This tampering might occur in-house or earlier, while in transit through the manufacturing supply chain.
6. D. Attackers can leverage misconfigured and insecure cloud deployments including unpatched applications, operating systems, and storage buckets.
7. A. Threat intelligence refers to knowledge about an existing or emerging threat to assets, including networks and systems. Threat intelligence includes context, mechanisms, indicators of compromise (IoCs), implications, and actionable advice.
8. A. Open-source intelligence (OSINT) applies to offensive security (ethical hacking/penetration testing) and defensive security. In offensive security, OSINT enables you to leverage public information from DNS records, social media sites, websites, search engines, and other sources for reconnaissance—in other words, to obtain information about a targeted individual or an organization. When it comes to threat intelligence, OSINT refers to public and free sources of threat intelligence.
9. D. Vendor websites, threat feeds, and vulnerability feeds could all be used for threat and vulnerability research.
10. C. The MITRE ATT&CK framework (https://attack.mitre.org) is a collection of different matrices of tactics and techniques. InfraGard is a collaborative effort between the FBI and the private sector. The Common Vulnerability and Exposure (CVE) is a standard to identify vulnerabilities created and maintained by MITRE. The Common Weakness Enumeration (CWE) is a standard to identify the weaknesses (root cause) of security vulnerabilities. CWE was also created by MITRE.

Review Questions

1. MITRE PRE-ATT&CK
2. STIX

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

3. TAXII
4. National Vulnerability Database (NVD)
5. Wireless
6. Authorized or ethical
7. Indicators of compromise (IoCs)
8. Advanced persistent threat (APT)
9. Shadow IT

Chapter 6

Do I Know This Already?

1. A. Infrastructure as a service (IaaS) is a type of cloud service that offers computer networking, storage, load balancing, routing, and VM hosting. Platform as a service (PaaS) provides various software solutions to organizations, especially the capability to develop applications in a virtual environment without the cost or administration of a physical platform. Software as a service (SaaS) is a cloud service model where the cloud provider offers the complete infrastructure and the application. Examples of SaaS include Gmail, Office 365, Webex, Zoom, Dropbox, Google Drive, and many other applications you use every day.
2. D. Encryption, authentication methods, and identity management are all security concerns in cloud deployments and environments.
3. B. A zero-day vulnerability is a type of vulnerability that is disclosed by an individual or exploited by an attacker before the creator of the software can create a patch to fix the underlying issue. Attacks leveraging zero-day vulnerabilities can cause damage even after the creator knows of the vulnerability because it may take time to release a patch to prevent the attacks and fix damage caused by them.
4. D. Default settings and passwords, weak encryption, and open permissions are examples of the most prevalent types of weak configurations that can be leveraged by an attacker to perform malicious activities and compromise systems.
5. D. Protocols such as Telnet, FTP (without encryption), and HTTP without encryption should be avoided at all times because they are considered unsecure.
6. D. Vendor management, system integration, lack of vendor support, supply chain, and outsourced code development should all be assessed when performing an analysis of third-party risks.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

7. A. A security update is a broadly released fix for a product-specific security-related vulnerability or group of vulnerability. Security vulnerabilities are rated based on their severity, which is indicated in the Microsoft Security Bulletin as critical, important, moderate, or low.
8. D. Legacy platforms and products that have passed the end of support date are often affected by unfixed security vulnerabilities and do not have modern security features. When a device is past the last day of support, vendors will not investigate or patch security vulnerabilities in those devices.
9. D. A security breach could have direct financial impact to a corporation (such as fines and lawsuits). The brand and reputation of a company can also be damaged by major cybersecurity incidents and breaches. Cybersecurity incidents can also lead to outages and availability loss.
10. D. Attackers can leverage different types of obfuscation and evasion techniques to go undetected (including encoding of data, tunneling, and encryption).

Review Questions

1. Cloud access security broker (CASB)
2. Zero-day vulnerability
3. Community cloud
4. Remote code execution (RCE)
5. SMTP with TLS encryption

Chapter 7

Do I Know This Already?

1. B. Threat hunting is the act of proactively and iteratively looking for threats in your organization.
2. A. The MITRE ATT&CK is a collection of matrices that outline adversary tactics, techniques, and procedures (TTPs) that modern attackers use.
3. A. Most of the vulnerabilities disclosed to the public are assigned Common Vulnerability and Exposure (CVE) identifiers. CVE is a standard created by MITRE (www.mitre.org) that provides a mechanism to assign an identifier to vulnerabilities so that you can correlate the reports of those vulnerabilities among sites, tools, and feeds.
4. C. A false positive is a broad term that describes a situation in which a security device triggers an alarm, but no malicious activity or actual attack is

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

taking place. In other words, false positives are false alarms, and they are also called benign triggers.

5. A. A true positive is a successful identification of a security attack or a malicious event. A true negative occurs when the intrusion detection device identifies an activity as acceptable behavior and the activity is actually acceptable. False positives are false alarms, and false negative is the term used to describe a network intrusion device's inability to detect true security events under certain circumstances—in other words, a malicious activity that is not detected by the security device.
6. A. Vulnerability scanners can often log in to the targeted system to perform deep analysis of the operating system, running applications, and security misconfigurations. This technique is called a credentialed scan.
7. D. SIEMs can provide log collection, normalization, and correlation.
8. A. NetFlow is a technology invented by Cisco to collect network metadata about all the different "flows" of traffic on your network.
9. A. Security Orchestration, Automation, and Response (SOAR) systems extend beyond traditional SIEMs to allow organizations to collect security threat data and alerts from multiple sources and to perform many different automated response capabilities.
10. D. Unlike traditional SIEM platforms, SOAR solutions can also be used for threat and vulnerability management, security incident response, and security operations automation (including playbook and runbook automation, as well as orchestration of multiple SOC tools).

Review Questions

1. A web application vulnerability scanner
2. Security advisories and bulletins
3. CVE Numbering Authorities (CNAs)
4. National Vulnerability Database (NVD)
5. Three
6. Medium
7. Threat hunting

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Chapter 8

Do I Know This Already?

1. D. Ethical hacking, pen testing, and penetration testing are all terms used to define the process of finding vulnerabilities and mimicking what an attacker could do against your systems and networks. Penetration testing is done after obtaining permission from the system or network owner.

2. A. In the known environment pen testing type, the pen tester starts out with a significant amount of information about the organization and its infrastructure. The tester would normally be provided network diagrams, IP addresses, configurations, and a set of user credentials, for example. If the scope includes an application assessment, the tester might also be provided the source code of the target application. The idea of this type of test is to identify as many security holes as possible.

3. D. The pre-engagement tasks include items such as contract negotiations, the statement of work (SOW), scoping, and the rules of engagement.

4. D. The rules of engagement document typically includes the testing timeline, location of the testing, time window of the testing, preferred method of communication, the security controls that could potentially detect or prevent testing, IP addresses or networks from which testing will originate, and the scope of the engagement.

5. A. *Open-source intelligence (OSINT) gathering* is the term used when a penetration tester uses public records to perform passive reconnaissance.

6. D. Active reconnaissance is carried out mostly by using network and vulnerability scanners. Nmap is an open-source network and port scanner. Nessus is a vulnerability scanner sold by Tenable. Nikto is an open-source web application vulnerability scanner.

7. B. The term *war driving* is used because the attacker can just drive around and get a huge amount of information over a very short period of time. A similar concept is war flying. In war flying an attacker can fly drones or similar devices to obtain information about a wireless network or even collect pictures or videos of facilities in some cases.

8. A. The blue team identifies the defenders of the organizations. Blue teams typically include the computer security incident response team (CSIRT) and information security (InfoSec) teams.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

9. C. Purple teams integrate the defensive capabilities of a blue team with the adversarial techniques used by the red team. Often the purple team is not a separate team, but a solid dynamic between the blue and red teams.
10. A. White teams are individuals who are focused on governance, management, risk assessment, and compliance.

Review Questions

1. Persistence
2. Privilege escalation
3. Partially known environment
4. Bug bounties
5. Passive

Chapter 9

Do I Know This Already?

1. A. Configuration management is an ongoing process created with the goal of maintaining computer systems, servers, network infrastructure, and software in a desired, consistent state. One of the primary goals of configuration management is to ensure that your infrastructure performs as it's expected to as changes are made over time.
2. C. After a minimum desired state of security is defined, baselines should be taken to assess the current security state of computers, servers, network devices, and the network in general. Baseline configurations should be properly documented and reviewed to include a set of specifications for information systems or configuration items within those systems. Baseline configurations are used by security professionals, along with network and system administrators, as a basis for future deployments, releases, or changes to information systems and applications.
3. A. You should make sure that your organization has appropriate naming conventions for describing IT infrastructure, applications, and users. Appropriate naming conventions are used to avoid conflicts and to be able to correlate data among disparate systems.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

4. B and C. The General Data Protection Regulation (GDPR) is a regulation in the European Union and the European Economic Area focused on data protection and privacy. Another example is the California Consumer Privacy Act (CCPA). These regulations give consumers the right to know what personal information is being collected by companies, government, and any other organizations.

5. D. Data loss prevention (DLP) systems can be software or hardware-based solutions and are categorized in three general types: endpoint, network, and storage DLP systems.

6. D. You should always encrypt data at rest, in use, and in motion in order to protect sensitive data.

7. C. Hashes are used in digital signatures, in file and message authentication, and as a way to protect the integrity of sensitive data—for example, data entered into databases or perhaps entire hard drives. A hash is generated through the use of a hash function to verify the integrity of the file or message, most commonly after transit over a network.

8. A. A hot site is a near duplicate of the original site of the organization that can be up and running within minutes (in some cases longer). Computers and phones are installed and ready to go, a simulated version of the server room stands ready, and the vast majority of the data is replicated to the site on a regular basis in the event that the original site is not accessible to users for whatever reason.

9. D. A cold site has tables, chairs, bathrooms, and possibly some technical setup—for example, basic phone, data, and electric lines. Otherwise, configuration of computers and data restoration are necessary before the site can be properly utilized. This type of site is used only if a company can handle the stress of being nonproductive for a week or more.

10. B. A honeypot is generally a single computer but could also be a file, group of files, or an area of unused IP address space, whereas a honeynet is a group of computers, servers, or an area of a network; a honeynet is used when a single honeypot is not sufficient. Either way, the individual computers, or group of servers, will usually not house any important company information. Various analysis tools are implemented to study the attacker; these tools, along with a centralized group of honeypots (or a honeynet), are known collectively as a honeyfarm.

11. D. Honeyfiles are used as bait files intended to lure adversaries to access them and then send alarms to a security analyst to potentially learn the tactics and techniques used by the attacker.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

12. A. In a DNS sinkhole you configure one or more DNS servers to provide false results to attackers and redirect them to areas in the network where you can observe their tactics and techniques. DNS sinkholes have been used to contain different types of malware such as the infamous WannaCry ransomware and to disrupt certain malicious DNS operations in denial-of-service (DoS) and other attacks.

Review Questions

1. Warm site
2. Honeyfiles
3. Digital rights management (DRM)
4. Data in use/processing
5. Tokenization

Chapter 10

Do I Know This Already?

1. A. IaaS is a service that offers computer networking, storage, load balancing, routing, and VM hosting. More and more organizations are seeing the benefits of offloading some of their networking infrastructure to the cloud.
2. A. A community cloud is a mix of public and private cloud deployments where multiple organizations can share the public portion.
3. A. Google Drive, Office 365, and Dropbox are examples of the software as a service (SaaS) cloud service model.
4. C. A managed security service provider (MSSP) provides services to manage your security devices and can also help monitor and respond to security incidents.
5. D. Managed service providers (MSPs) can deliver network, application, system, and management services using a pay-as-you-go model. An MSP is an organization that can manage your network infrastructure, servers, and in some cases your security devices. Companies that provide services to manage your security devices and can also help monitor and respond to security incidents are called managed security service providers (MSSPs).
6. C. The term *edge computing* describes an ecosystem of resources and applications in new network services (including 5G and IoT). One of the main benefits is to provide greater network speeds, low latency, and computational power near the user.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

7. A. Thin clients are computer systems that run from resources stored on a central server or from the cloud instead of a local (on-premises) system. When you use a thin client, you connect remotely to a server-based computing environment where the applications, sensitive data, and memory are stored.
8. D. Docker Swarm, Apache Mesos, and Kubernetes are technologies and solutions to manage, deploy, and orchestrate containers.
9. A. VM sprawl (otherwise known as virtualization sprawl) occurs when an organization can no longer effectively control and manage all the VMs on a network or in the cloud.
10. A. In a VM escape attack, the guest VM breaks out of its isolated environment and attacks the hypervisor or compromises other VMs hosted and controlled by the hypervisor.

Review Questions

1. Transit gateway
2. Resource policies
3. Cloud services integration
4. Serverless architecture
5. SDN

Chapter 11

Do I Know This Already?

1. A. The traditional software development methodology is the waterfall model, which is a software and hardware development and project management methodology that has at least five to seven phases that follow in strict linear order. Each phase cannot start until the previous phase has been completed.
2. B. You can integrate with log aggregation tools to maintain and analyze logs of every element that goes into the provisioning. This allows you to respond quickly and deprovision the application in the event that something went wrong. When you go back, you can check the logs and accurately find and fix the root cause of the error.
3. D. Unit testing, integration testing, and identifying a code integrity manager can help software (code) integrity.
4. C. Threat modeling enables you to prioritize threats to an application based on their potential impact. This modeling process includes identifying assets to the system or application, uncovering vulnerabilities, identifying threats,

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

documenting threats, and rating those threats according to their potential impact. The more risk, the higher the rating. Threat modeling is often incorporated into the software development lifecycle (SDLC) during the design, testing, and deployment phases.

5. D. Input validation, principle of least privilege, and failing securely are all important security principles that should be incorporated into the SDLC.

6. D. Compile time refers to the duration of time during which the statements written in any programming language are checked for errors. Compile-time errors might include syntax errors in the code and type-checking errors. A programmer can check these errors without actually running the program and instead check it in the compile stage when it is converted into machine code.

7. D. One of the most popular OWASP projects is the Top 10 Web Application Security Risks. You can find the latest Top 10 Web Application Security Risks at https://owasp.org/www-project-top-ten. All of these answers are top web application security risks.

8. C. In an example of software diversity, a compiler is modified to generate variants of a binary (target application) that operates in the same way when processing benign input; however, it may operate in a different manner when given malicious input. This new aspect of software diversity is handled by generating variants of a program by building a binary with a diversifying compiler that can randomize the code layout, stack variables, and random allocations of heap objects at different locations in each variant.

9. A. Continuous integration (CI) is a software development practice in which programmers merge code changes in a central repository multiple times a day. Continuous delivery (CD) sits on top of CI and provides a way for automating the entire software release process. When you adopt CI/CD methodologies, each change in code should trigger an automated build-and-test sequence. This automation should also provide feedback to the programmers who made the change.

10. C. Elasticity is the ability of an underlying infrastructure to react to a sudden increase in demand by provisioning more resources in an automated way. Elasticity and scalability are often achieved by deploying technologies such as load balancers and by deploying applications and resources in multiple geographical locations (data centers around the world). Other technologies such as enabling concurrent processing (parallel processing) and automated container deployments (that is, using Kubernetes) allow organizations to auto-scale.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Review Questions

1. Static analysis
2. Software integrity measurement
3. Development
4. Staging
5. Agile

Chapter 12

Do I Know This Already?

1. D. Microsoft Active Directory (AD) allows administrators to organize elements of a network, such as users, computers, and devices, into a hierarchical containment structure.
2. D. Users authenticate against a directory service to ensure that it is highly available at all times. A best practice requires that it be distributed in multiple locations.
3. C. In biometrics, facial recognition is the most common and least accurate way to identify a user; it has higher false rejection and higher false acceptance rates than other biometric security methods.
4. B. Speaker verification is a 1:1 match where one speaker's voice is matched to one template, also called a voice print or voice model.
5. B. You must have this physical item with you—something you have such as a crypto card, token, or key fob that is used as a method to authenticate you as a user.
6. A. Physical movements such as the way you walk, typing patterns, and mouse movements are examples of a personality trait, behavior, or observable biometric that can be used to authenticate you as a person.
7. A. Authentication, authorization, and accounting (AAA) is a framework for intelligently controlling access to computer resources, enforcing policies, and auditing usage. These processes working in concert are important for effective network management and security.
8. A. Authentication provides a method of identifying a user, typically by having the user enter a valid username and password before access to the network is granted. Authentication is based on each user having a unique set of login credentials for gaining network access.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

9. C. Enterprises that elect to use a cloud computing model need to pay only for the resources that they use, with none of the maintenance and upkeep costs. The price adjusts up or down depending on how much is consumed.

10. B. On-premises authentication is preferred. Many companies these days operate under some form of regulatory control, regardless of the industry. The most common one is the Health Insurance Portability and Accountability Act (HIPAA) for private health information, but there are other government and industry regulations. For companies that are subject to such regulations, it is imperative that they remain compliant and know where their data is at all times.

Review Questions

1. Directory services
2. Lookup
3. Federation
4. Vein or vein authentication
5. Gait analysis
6. The crossover error rate (CER) describes the point where the false rejection rate (FRR) and false acceptance rate (FAR) are equal. The crossover error rate describes the overall accuracy of a biometric system.

Chapter 13

Do I Know This Already?

1. A. Geographic dispersal of computing and data assets if a disaster, natural or person-made, occurs ensures the company can continue to function. In such cases, the company relies on infrastructure in another city, state, or country to be available.

2. A. RAID 5 writes the parity over all disks, making it possible to continue to run even if one disk fails. Recovery is done by removing the failed drive and executing the recovery process, where the drive is rebuilt and added to the array once completed.

3. A. RAID 0 is known as disk striping; it is the process of dividing data into blocks and spreading the data blocks across multiple storage devices, such as hard disks or solid-state drives (SSDs). In a Redundant Array of Independent Disks (RAID) group, there is no parity. The more hard drives in a RAID 0 array, the higher probability of array failure.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

4. B. By deploying NIC teaming on your server, you can maintain a connection to multiple physical switches and use a single IP address. Load balancing becomes readily available, fault tolerance becomes instant instead of waiting for DNS records to time out or update, and management becomes simpler.

5. B. Redundant power is a critical component in high-availability systems. In the simplest solution, two power supplies can drive a load through bused (vertical line shared among equipment) or N+1 configuration, where two power supply outputs are load-shared together or have one active and one or more in standby mode.

6. B. SAN-connected servers contain special fiber interface cards called host bus adapters (HBAs). They are configured as pairs, typically called HBA1 and HBA2. The fiber is then connected to a pair of SAN network switches.

7. A. Azure Active Directory is a fully managed multi-tenant cloud-based offering from Microsoft that offers identity and access capabilities for applications running in an on-premises environment. It is not a replacement for on-premises Active Directory Services but could be, or it can be used with it to extend on-prem directory services and sync the directories to cloud applications.

8. D. Cloud-based backups are a model of data storage in which the data can be accessed, managed, and stored in a remote cloud server via the Internet. Cloud backups are maintained and supported by a cloud storage provider responsible for keeping the user's data available and accessible at any time.

9. C. Backups to a cloud services backup (IaaS) provider can be slower than on-premises backups where they are connected at gigabit speeds. Depending on the amount of data being backed up, restoring data from them could take many hours or even days to complete.

10. D. A highly available system should be able to quickly recover from any sort of failure state to minimize interruptions for the end user. Best practice for achieving high availability is to eliminate single points of failure or any node that would impact the system as a whole if it becomes dysfunctional. The highest level of uptime is considered "five nines," or 99.999 percent, which refers to a standard of reliability. Five nines is equivalent to downtime of only 5 minutes and 15 seconds per year (1 minute and 18 seconds in a quarter, or 26 seconds monthly). These are very high standards to meet.

11. D. Without the network, systems will be unable to communicate with each other. When building the restore order for your organization, do not forget this critical step, and make sure you add it to your desktop exercises.

12. A, B, C. These are all steps that can be taken to enhance an organization's resilience and provide fault tolerance and diversity.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Review Questions

1. Geographical dispersal is the practice of placing valuable data assets around the city, state, country, or world to provide an extra level of protection from attacks, mistakes, and disasters.
2. In the simplest of terms, disk redundancy is a system's ability to write data to two or more disks at the same time. Having the same data stored on separate disks enables you to recover the data in the event of a disk failure.
3. An uninterruptible power supply or uninterruptible power source (UPS) is an electrical device that provides emergency power to a load when the input power source or mains power fails. Generally, UPSs are battery based—a bank of batteries and circuits that provide power during main power failure.
4. Data replication via a SAN is the most common method of replication. Replicating data from one data center to another via dual SANs allows you to replicate large volumes of data quickly using SAN technology.
5. Reverting to known state is returning the system to a state prior to a specific moment in time or state of existence.

Chapter 14

Do I Know This Already?

1. A. A field-programmable gate array (FPGA) is an integrated circuit designed to be configured by a customer or designer after manufacturing—hence the term *field-programmable*.
2. A. SCADA systems are capable of managing parts inventories for just-in-time manufacturing, regulate industrial automation and robots, and monitor processes and quality control.
3. A. A SCADA programmable logic controller (PLC) is an industrial computer control system that continuously monitors the state of input devices and makes decisions based on a custom program to control the state of output devices.
4. B. Smart watches can monitor your pulse, heart, blood pressure, exercises, calories, and sleep patterns to help you become more fit.
5. A. Lighting and air conditioning controls are part of building automation and can be controlled to help reduce the overexpense of energy use by properly implementing sensors to know when people are in the building.
6. D. Users should utilize the strongest encryption method available on devices, always create and rotate very complex passwords, set up a continuous method

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

to log and audit access to the devices, and finally, ensure use of the latest manufacturer updates.

7. D. Today's modern surveillance systems are capable of performing facial recognition, monitoring data center access, and if exploited, pivoting to and compromising other systems in the network.

8. D. Zigbee creates flexibility for developers and end users while delivering stellar interoperability. It was created on the IEEE's 802.15.4 standard, using the 2.4-GHz band and a self-healing true mesh network, and has a defined rate of 250 kbps. It is best suited for intermittent data transmissions from a sensor or input device.

9. C. Cryptography constraints for building secure embedded systems hardware and software have to do with the amount of code required to implement a secure algorithm and the processing power required to crunch the number.

10. C and D. A system on a chip (SoC) is essentially an integrated circuit or an IC that takes a single platform and integrates an entire electronic or computer system onto it.

11. A. Nucleus RTOS from Mentor is a real-time operating system (RTOS) software component that rapidly switches between tasks, giving the impression that multiple programs are being executed at the same time on a single processing core. In fact, the processing core can execute only one program at any one time.

12. B. Embedded computers used in unmanned system applications are often characterized by their low SWaP-C (size, weight, power, and cost) profiles, small form factor (SFF), and rugged operating ranges, which are vital components for unmanned aerial vehicles (UAVs),

Review Questions

1. Arduino devices are hardware and software combined into an extremely flexible platform; they can read inputs such as a light on a sensor or a button press. The Arduino software is easy for beginners to use yet flexible enough for advanced users. It runs on Mac, Windows, and Linux.

2. Field-programmable gate arrays (FPGAs) are integrated circuits designed to be configured by a customer or designer after manufacturing—hence the term *field programmable*.

3. In the manufacturing process, control systems can help with the reduction of product errors and discards, due to earlier problem detection and remedies. These systems improve productivity, maximizing the effectiveness of machine uptime.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

4. Cybersecurity and attacks on the platform are the biggest problem with IoT devices being developed and sold. Cybersecurity must be designed into IoT devices from the ground up and at all points in the ecosystem to prevent vulnerabilities in one part from jeopardizing the security of the entire system.

5. The CAN bus system enables each ECU to communicate with all other ECUs, without complex dedicated wiring. An ECU can prepare and broadcast information (that is, sensor data) via the CAN bus, consisting of two wires—CAN low and CAN high. The broadcasted data is accepted by all other ECUs on the CAN network.

Chapter 15

Do I Know This Already?

1. A. An access control vestibule, formerly known as a mantrap, allows security guards to see visitors before they are allowed through the second door. Guards can use cameras and voice communication to ascertain identity.

2. C. Vestibules are an excellent access control addition. Entries with panels built from prefabricated composite or metal are used as a way for companies to control the heat and airflow in their facilities.

3. B. Optical detectors convert incoming optical energy into electrical signals. The two main types of optical detectors are photon detectors and thermal detectors. Photon detectors produce one electron for each incoming photon of optical energy. The electron is then detected by the electronic circuitry.

4. D. Robot sentries act as 24/7/365 guards, continuously monitoring and alerting on differentials. Robot sentries report anything out of the ordinary to the appropriate personnel, who then can take additional action.

5. D. The reception desk plays a front-line role in the physical security program of a company. Receptionists do this by creating a buffer between the corporate offices, employees, and contractors. Visitors are unable to pass until the employees meeting them come and pick them up. This reduces loitering and inquisitive visitors.

6. D. Biometric locks provide a unique way of making sure people are who they say they are by monitoring or matching human characteristics such as a fingerprint, retina, or voice prior to unlocking and granting access.

7. A. A skeleton key normally work with warded or lever locks. With a warded lock, a skeleton key lacks interior notches to interfere with or correspond with the wards, or obstructions, thereby allowing it to open the lock.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

8. C. A proximity reader is capable of reading a prox card that is within a few millimeters of the reader/pad; it does this through induction.
9. C. A vault can consist of an entire room or even multiple rooms. Vaults allow everything inside to be protected through multiple layers of security measures, including guards, alarms, cameras, locks, gates, and secure doors.
10. D. Pulverizing grinds devices down to bare-metal scraps. There is nothing left that would allow recovery, unlike shredding and degaussing.

Review Questions

1. An access badge is a credential used to gain entry to an area having automated readers for access control entry points.
2. Appropriately placed signage provides direction and guidance for staff and visitors; it also provides clear expectations and the repercussions for failure to abide by those rules.
3. It enables corporate, industrial, and data centers to blend into their environment. When you surround the premises with trees, bushes, and vegetation and implement low-profile security measures around the perimeter, the building becomes one with the area. This ensures it does not stick out and become a highly visible target.
4. One of the two people is there as an observer; this person monitors the person performing the work and ensures that person is performing work exactly as described in the change request and can also question any variance. The monitor typically reports any unusual or suspicious activity immediately to security or the guards' office.
5. A proximity reader or prox reader, typically an RFID reader, reads a card by placing it near (within proximity of) the reader. The reader sends energy in the form of a field to the card, powering up the card, which enables the reader to read the information stored on the prox card. Prox cards are used as part of an access control system.

Chapter 16

Do I Know This Already?

1. A. A key generation algorithm that selects a private key uniformly at random from a set of possible private keys is one of three digital signature scheme algorithms. This algorithm outputs the private key and a corresponding public key.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

2. B. Cryptographic hash functions have many information-security applications, notably in digital signatures, message authentication codes (MACs), and other forms of authentication.
3. D. Quantum cryptography, or quantum key distribution (QKD), uses a series of photons (light particles) to transmit data from one location to another over a fiber-optic cable. By comparing measurements of the properties of a fraction of these photons, the two endpoints can determine what the key is and if it is safe to use.
4. B. Photons travel to a receiver, which uses two beam splitters (horizontal/vertical and diagonal) to "read" the polarization of each photon.
5. A. A blockchain is essentially a digital ledger of transactions that are duplicated and distributed across the entire network of computer systems on the blockchain.
6. B. Cipher suites usually contain a set of algorithms that include a key exchange algorithm, bulk encryption algorithm, and message authentication code (MAC) algorithm. The key exchange algorithm is used to exchange a key between two devices.
7. D. Steghide, Foremost, Xiao, Stegais, and Concealment are tools that can be used to conceal data in images and audio files.
8. A. The most common use cases are mobile devices and portable systems. Because of the low-power draw requirements, you may use smaller symmetric key sizes and elliptic-curve asymmetric encryption.
9. C. Time is used in generating keys used in cryptography. They have been mostly replaced with more randomness.
10. D. Key stretching runs a password through an algorithm to produce an enhanced key.
11. C and D. Bcrypt and Password-Based Key Derivation Function 2 (PBKDF2) are key derivation functions (KDFs) that are primarily used for key stretching, which provides a means to "stretch" a key or password, making an existing key or password stronger and protecting against brute-force attacks.

Review Questions

1. Digital signatures employ asymmetric cryptography. In many instances, they provide a layer of validation and security to messages sent through a nonsecure channel. Properly implemented, a digital signature gives the receiver reason to believe the message was sent by the claimed sender.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

2. Key stretching techniques are used to make a possibly weak key, typically a password or passphrase, more secure against brute-force attacks by increasing the resources (time and possibly space) needed to test each possible key.
3. Salts defend against a precomputed hash attack. Because salts are different in each case, they also protect commonly used passwords, or those users who use the same password on several sites, by making all salted hash instances for the same password different from each other.
4. Block ciphers are an encryption method that applies a deterministic algorithm along with a symmetric key to encrypt a block of text instead of encrypting one bit at a time as in stream ciphers.
5. Asymmetric encryption is also known as public key cryptography; asymmetric encryption uses two keys to encrypt plaintext. Secret keys are exchanged over the network. This type of encryption ensures that malicious persons do not misuse the keys. It is important to note that anyone with the secret key can decrypt the message, and this is why asymmetric encryption uses two related keys to boost security.
6. In an asymmetric key system, each user has a pair of keys: a private key and a public key. To send an encrypted message, you must encrypt the message with the recipient's public key. The recipient then decrypts the message with his or her private key. The easiest thing to remember is that public keys encrypt and private keys decrypt.
7. Ephemeral describes something of a temporary or short duration. Ephemeral keys are designed to be used for a single transaction or session. The term *ephemeral* is increasingly used in computer technology.

Chapter 17

Do I Know This Already?

1. D. DNSSEC is used in securing the chain of trust that exists between the Domain Name System (DNS) records that are stored at each domain level.
2. B. You can enable LDAPS by installing a properly formatted certificate from a certificate authority (CA) according to the guidelines.
3. B. SSH uses asymmetric (public key) RSA cryptography for both connection and authentication.
4. B. S/MIME is based on asymmetric cryptography that uses a pair of mathematically related keys to operate: a public key and a private key.
5. C. AH is an optional packet header used to guarantee connectionless integrity and data origin authentication for IP packets.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

6. C. IMAPS downloads a message only when you click on it, and attachments aren't automatically downloaded. IMAPS operates on port 993 (SSL/TLS).
7. D. HTTPS uses an encryption protocol to encrypt communications. The protocol is called Transport Layer Security (TLS).
8. D. Enterprises looking to deploy time synchronization should utilize three public servers, set up a local internal NTP server that is used for all internal hosts as a reference timekeeper, and only have the internal NTP server make requests to the public servers. You should ensure that you are standardizing on UTC time across all systems; it will make researching attacks and issues more relevant.
9. A. A virtual private network (VPN) provides privacy and security to users by creating a private network connection across a public network connection. VPNs are used to access company networks protected by firewalls that deny inbound and outbound access to/from systems on the network.
10. A, B, C, D. All the responses are correct. You should use caching-only DNS servers, use DNS forwarders, use DNS advertisers and resolvers, protect DNS from cache pollution, enable DDNS for secure connections only, disable zone transfers, and use firewalls to control communication to and from the DNS servers.

Review Questions

1. A cryptographic protocol or encryption protocol is an abstract of a protocol that performs a security-related function and applies cryptographic methods, often as sequences of cryptographic primitives.
2. SSH uses encryption to ensure secure transfer of information between the host and the client. Host refers to the remote server you are trying to access, and the client is the computer you are using to access the host.
3. S/MIME is based on asymmetric cryptography, which uses a pair of mathematically related keys to operate: a public key and a private key.
4. SRTP and SRTCP use the Advanced Encryption Standard (AES) as the default cipher.
5. You can enable LDAPS by installing a properly formatted certificate from a certificate authority (CA) according to the guidelines. LDAPS over SSL/TLS uses TCP port 636.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Chapter 18

Do I Know This Already?

1. A. According a report published by Cybersecurity Ventures in May 2019, damages from ransomware cost businesses an astonishing $11 billion in lost revenue, productivity, and remediation.
2. B. Antimalware software that uses behavior-based malware detection can detect previously unknown threats by identifying malware based on characteristics and behaviors.
3. C. One of the primary functions of endpoint detection and response (EDR) is providing forensics and analysis tools to research identified threats and search for suspicious activities (similar to threat hunting).
4. A. Data loss prevention (DLP) software and tools monitor and control endpoint activities, filter data streams on corporate networks, and monitor data in the cloud to protect data at rest and in motion.
5. B. The newer specification addresses several limitations of the BIOS, including restrictions on hard disk partition size and the amount of time the BIOS takes to perform its tasks.
6. D. Database tokenization is the process of turning sensitive data into nonsensitive data called *tokens* that can be used in a database or internal system without bringing it into scope.
7. D. The secure session cookies store information about a user session after the user logs in to an application. This information is very sensitive because an attacker can use a session cookie to impersonate the victim.
8. A. Patch management is a process that helps acquire, test, and install multiple patches (code changes) on existing applications and software tools on a computer, enabling systems to stay updated on existing patches and determining which patches are the appropriate ones.
9. C. When a self-encrypting drive (SED) is installed into a mixed-disk configuration or a configuration containing unencrypted drives, it operates as an unencrypted disk. Likewise, a pool consisting of all SEDs might replicate to a pool with only a few SEDs or no SEDs at all.
10. C. Code from the outside needs to be validated prior to running it on a secure CPU. This tamper resistance can be implemented in many ways—for example, using a dedicated ROM that can only be accessed by the hardware root of trust.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

11. D. The TPM uses a unique key to digitally sign the log recorded by the UEFI.

12. A. Sandboxing is a strategy that isolates a test environment for applications. It provides an extra layer of security that prevents malware or harmful applications from negatively affecting your system.

Review Questions

1. Antimalware software uses signature-based detection, behavior-based detection, and sandboxing.

2. Boot integrity refers to using a secure method to boot a system and verify the integrity of the operating system and loading mechanism. Boot integrity represents the first step toward achieving a trusted infrastructure.

3. In boot attestation, software integrity measurements are immediately committed to during boot, thus relaxing the traditional requirement for secure storage.

4. Full-disk encryption (FDE) is a cryptographic method that applies encryption to the entire hard drive, including data, files, operating system, and software programs. FDE encryption places an exterior guard on the internal contents of the device.

5. Self-encrypting drives (SEDs) are disk drives that use an encryption key to secure the data stored on the disk. This encryption protects the data and array from data theft when a drive is removed from the array. Because SED operates across all disks in an array at once, the drive must be configured as an SED when introduced to the array.

Chapter 19

Do I Know This Already?

1. A. If you migrate some of these low-resource servers to a virtual environment, you could end up spending more on licensing but less on hardware, due to the nature of virtualization. In fact, the goal is to have the gains of hardware savings outweigh the losses of licensing. Load balancing and clustering deal with an operating system utilizing the hardware of multiple servers. This will not be the case when you go virtual, nor would it have been the case anyway, because clustering and load balancing are used in environments where the server is very resource-intensive. Baselining, unfortunately, will remain the same; you should analyze all of your servers regularly, whether they are physical

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

or virtual. These particular servers should not encounter latency or lowered throughput because they are low-resource servers in the first place. If, however, you considered placing a Windows Server that supports 5000 users into a virtual environment, you should definitely expect latency.

2. B. You can defend against pivoting by providing proper access control, network segmentation, DNS security, reputation security, and proper patch management.

3. C. A remote-access VPN is typically used for client access to a headend device, which connects them to the corporate network. Most remote access VPNs use IPsec or SSL/TLS connections.

4. D. One specific type of DDoS is the DNS amplification attack. Amplification attacks generate a high volume of packets ultimately intended to flood a target website. In the case of a DNS amplification attack, the attacker initiates DNS requests with a spoofed source IP address. The attacker relies on reflection; responses are not sent back to the attacker but are instead sent "back" to the victim server. Because the DNS response is larger than the DNS request (usually), it amplifies the amount of data being passed to the victim. An attacker can use a small number of systems with little bandwidth to create a sizable attack. However, a DNS amplification attack can also be accomplished with the aid of a botnet, which has proven to be devastating to sections of the Internet during the period when the attack was carried out.

5. D. Domain hijacking is a type of hijacking attack in which the attacker changes the registration of a domain name without the permission of the original owner/registrant. One of the most common methods used to perform a domain hijacking is using social engineering.

6. B. Some companies (such as Cisco) offer hardware-based NAC solutions, whereas other organizations offer paid software-based NAC solutions and free ones such as PacketFence (https://packetfence.org), which is open source. The IEEE 802.1X standard, known as port-based network access control, or PNAC, is a basic form of NAC that enables the establishment of authenticated point-to-point connections, but NAC has grown to include software; 802.1X is now considered a subset of NAC.

7. B. As administrator, you might need to take an alternate path to manage network devices. In this case, you might require out-of-band management. This is common for devices that do not have a direct network connection, such as UPSs, PBX systems, and environmental controls.

8. D. Port security is a security feature present in most routers and switches, and it is used to provide access control by restricting the Media Access Control (MAC) addresses that can be connected to a given port. This differs from a

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

MAC access list because it works only on the source MAC address without matching the MAC destination.

9. D. IP proxy secures a network by keeping machines behind it anonymous; it does this through the use of NAT. For example, a basic four-port router can act as an IP proxy for the clients on the LAN it protects. An IP proxy can be the victim of many network attacks, especially DoS attacks. Regardless of whether the IP proxy is an appliance or a computer, it should be updated regularly, and its log files should be monitored periodically and audited according to organization policies.

10. D. A Layer 2 access control list (ACL) operates at the data link layer of the OSI model and implements filters based on Layer 2 information. An example of this type of access list is a MAC access list, which uses information about MAC addresses to create the filter.

11. D. There are different route manipulation attacks, but one of the most common is the Border Gateway Protocol (BGP) hijacking attack. BGP is a dynamic routing protocol used to route Internet traffic. An attacker can launch a BGP hijacking by configuring or compromising an edge router to announce prefixes that have not been assigned to his or her organization. If the malicious announcement contains a route that is more specific than the legitimate advertisement or presents a shorter path, the victim's traffic may be redirected to the attacker. In the past, threat actors have leveraged unused prefixes for BGP hijacking to avoid attention from the legitimate user or organization.

12. D. As administrator, you can use QoS capabilities to control application prioritization. Protocol discovery features in Cisco AVC show the mix of applications currently running on the network. This information helps you define QoS classes and policies, such as how much bandwidth to provide to mission-critical applications and how to determine which protocols should be policed. Per protocol bidirectional statistics are available, such as packet and byte counts, as well as bit rates. After you classify the network traffic, you can apply class-based weighted fair queuing (CBWFQ) for guaranteed bandwidth.

13. D. An anycast address is assigned to a group of interfaces on multiple nodes. Packets are delivered to the "first" interface only. Anycast is structured like unicast addresses.

14. A. To capture packets and measure throughput, you need a tap on the network before you can start monitoring. Most tools that collect throughput leverage a single point configured to provide raw data, such as pulling traffic from a switch or router. If the access point for the traffic is a switch, typically a network port is configured as a Switched Port Analyzer (SPAN) port, sometimes also called port mirroring or port spanning. The probe capturing data from a

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

SPAN port can be either a local probe or data from a SPAN port that is routed to a remote monitoring tool.

15. A. Baselining is the process of measuring changes in networking, hardware, software, applications, and so on. The process of documenting and accounting for changes in a baseline is known as baseline reporting. Baseline reporting enables you to identify the security posture of an application, system, or network. The security posture can be defined as the risk level to which a system, or other technology element, is exposed. Security posture assessments (SPAs) use baseline reporting and other analyses to discover vulnerabilities and weaknesses in systems.

16. C. File integrity is important when securing log files, and File Integrity Monitoring (FIM) helps you maintain this integrity. Encrypting the log files through the concept known as hashing is a good way to verify the integrity of the log files if they are moved and/or copied. You could also flat-out encrypt the entire contents of the file so that other users cannot view it. Integrity means that data has not been tampered with. Authorization is necessary before data can be modified in any way; this is done to protect the data's integrity.

Review Questions

1. Cost
2. Network segmentation
3. Clientless
4. DNS amplification
5. Network access control (NAC)
6. Out-of-Band
7. Port security
8. Proxy server
9. Layer 2
10. Border Gateway Protocol (BGP)
11. Class-based weighted fair queuing (CBWFQ)
12. Anycast

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Chapter 20

Do I Know This Already?

1. A. WPA3 includes a more robust authentication mechanism than WPA2. It also provides a higher level of encryption capabilities. It enables a very robust authentication based on passwords by utilizing a technology called Simultaneous Authentication of Equals (SAE). This innovation in Wi-Fi security replaces the preshared key (PSK). SAE helps protect against brute-force password attacks and offline dictionary attacks.

2. B. Counter-mode/CBC-MAC protocol or CCMP is based on the Advanced Encryption Standard (AES). It provides a stronger mechanism for securing privacy and integrity over Temporal Key Integrity Protocol (TKIP), which was previously used with WPA. An advantage to CCMP is that it utilizes 128-bit keys as well as a 48-bit initialization vector. This enhancement greatly reduces the possibility of replay attacks. One drawback to using CCMP over TKIP is that it requires additional processing power. That is why you will typically see it supported on newer hardware.

3. A. 802.1X is an IEEE standard that defines port-based network access control (PNAC). Not to be confused with 802.11X WLAN standards, 802.1X is a data link layer authentication technology used to connect hosts to a LAN or WLAN. 802.1X allows you to apply a security control that ties physical ports to end-device MAC addresses, and prevents additional devices from being connected to the network. It is a good way of implementing port security, much better than simply setting up MAC filtering.

4. D. Following are the three components to an 802.1X connection:

   - **Supplicant:** A software client running on a workstation. This is also known as an authentication agent.
   - **Authenticator:** A wireless access point or switch.
   - **Authentication server:** An authentication database, most likely a RADIUS server.

5. D. The supplicant is a software client running on a workstation. It is also known as an authentication agent.

6. D. The Protected Extensible Authentication Protocol (PEAP) uses MS-CHAPv2, which supports authentication via Microsoft Active Directory databases. It competes with EAP-TTLS and includes legacy password-based protocols. It creates a TLS tunnel by acquiring a public key infrastructure (PKI)

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

certificate from a server known as a certificate authority (CA). The TLS tunnel protects user authentication much like EAP-TTLS.

7. B. EAP-FAST uses a protected access credential instead of a certificate to achieve mutual authentication. FAST stands for Flexible Authentication via Secure Tunneling.

8. C. The preshared key (PSK) used to enable connectivity between wireless clients and the WAP is a complex passphrase. PSK is automatically used when you select WPA-Personal in the Security Mode section. The other option is WPA-Enterprise, which uses a RADIUS server. So, if you ever see the term *WPA2-PSK*, this means that the WAP is set up to use the WPA2 protocol with a preshared key, and not an external authentication method such as RADIUS.

9. B. Wi-Fi Protected Setup (WPS) is a security vulnerability. Created originally to give users easy connectivity to a wireless access point, later all major manufacturers suggested that it be disabled (if possible). In a nutshell, the problem with WPS was the eight-digit code. It effectively worked as two separate smaller codes that collectively could be broken by a brute-force attack within hours.

10. D. Strategic wireless access point (WAP) placement is vital. That is why it is essential to perform a site survey before deploying wireless equipment. A site survey is typically performed using Wi-Fi Analyzer tools to produce a heat map of all wireless activity in the area.

Review Questions

1. SAE
2. EAP-FAST
3. CCMP
4. Supplicant
5. Authenticator
6. Wi-Fi Analyzer
7. AES

Chapter 21

Do I Know This Already?

1. A. Near-field communication (NFC) has obvious benefits for contactless payment systems or any other non-contact-oriented communications between devices. However, for optimal security, you should use contact-oriented readers and cards.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

2. B. Geofencing is an excellent way to be alerted to users entering and exiting an organization's physical premises. It can provide security for wireless networks by defining the physical borders and allowing or disallowing access based on the physical location of the user, or more accurately, the user's computer or mobile device.
3. A. You should encrypt data communication between a device and the organization and enable full device encryption of data stored on the device or in removable storage.
4. D. One of the characteristics of an MDM solution is the use of over-the-air (OTA) device management updates. OTA historically refers to the deployment and configuration performed via a messaging service, such as Short Message Service (SMS), Multimedia Messaging Service (MMS), or Wireless Application Protocol (WAP). Now it's used to indicate remote configuration and deployment of mobile devices.
5. D. Insecure user configurations such as rooting and jailbreaking can be blocked from MDM, as can sideloading—the art of loading third-party apps from a location outside the official application store for that device. Note that sideloading can occur in several ways: by direct Internet connection (usually disabled by default), by connecting to a second mobile device via USB OTG (USB On-The-Go) or Bluetooth, by copying apps directly from a microSD card, or by tethering to a PC or Mac.
6. B. Bluejacking is the sending of unsolicited messages to Bluetooth-enabled devices such as mobile phones. You can stop bluejacking by setting the affected Bluetooth device to undiscoverable or by turning off Bluetooth altogether.
7. B. Encryption is one of the best ways to ensure that data is secured and that applications work properly without interference from potential attackers. However, you should consider whole device encryption, which encrypts the internal memory and any removable (SD) cards.
8. B. Companies may implement strategies such as choose your own device (CYOD), where employees select a device from a company-approved list, or corporate-owned, personally enabled (COPE), where the company supplies employees with phones that can also be used for personal activities.

Review Questions

1. Virtual desktop infrastructure (VDI)
2. SEAndroid
3. Geotagging

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

4. Application deny/block list
5. Bluejacking
6. USB On-The-Go (USB OTG)
7. Bluesnarfing
8. Sideloading

Chapter 22

Do I Know This Already?

1. A. Resource policies in cloud computing environments are meant to control access to a set of resources. A policy is deployed to manage the access to the resource itself.
2. B. In cloud computing environments, the management of things like API keys, passwords, and certificates is typically handled by some kind of secrets management tool. It provides a mechanism for managing the access and auditing the secrets used in the cloud environment.
3. A. In cloud computing environments, storage is referred to as buckets. The access to these buckets is controlled by an Identity and Access Management (IAM) policy.
4. D. In cloud computing environments, the concept of a public subnet is one that has a route to the Internet. A private subnet would be a subnet in the cloud environment that does not have a route to the Internet.
5. D. A cloud access security broker (CASB) is a tool that organizations utilize to control access to and use of cloud-based computing environments and resources. Many cloud-based tools are available for corporate and personal use. The flexible access nature of these tools makes them a threat to data leak prevention and the like. It is very easy to make the mistake of copying a file that contains sensitive data into the wrong folder and making it available to the world. This is the type of scenario that CASB solutions help mitigate.
6. B. The concept of a Secure Web Gateway (SWG) is top of mind these days. At the time of this writing, we are currently in a global pandemic that has forced millions to work from home. For many years, the solution for providing employees a way to work from home was simply a remote-access VPN back into the office. This solution allowed all of the traffic from the employees' computers to flow back through the corporate network, which would in turn traverse the same security controls that were in place if the employees were plugged into the corporate network. With the advent of cloud-based access to

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

applications and storage, this solution is no longer the most efficient way of securing remote workers. This is where the SWG comes into play. A Secure Web Gateway enables you to secure your remote workers' Internet access while not overloading the corporate Internet pipe. This approach is sometimes thought of as a cloud firewall. However, an SWG typically has many other protection mechanisms in place, including things like CASB. One example of an SWG is Cisco Umbrella.

7. B. Cloud-based firewalls can work at all layers of the OSI model. However, in many cloud computing environments, the firewall is used at the application layer to control access and mitigate threats to applications being hosted by the cloud environment.

8. B. A cloud native control is typically provided by the actual cloud-computing environment vendor. A non-cloud native control is provided by a third-party vendor. For instance, each cloud computing environment has security controls built into its platform. However, these controls might not be sufficient for all use cases. That is where third-party solutions come into play. Many companies out there today provide these third-party solutions to supplement those areas where cloud native controls are lacking. Of course, most virtual machine–based controls can be deployed in any cloud computing environment. However, for it to be native, it must go through various integration, testing, and certification efforts. Because each cloud computing environment is built differently and on different platforms, the controls must be adapted to work as efficiently as possible with the specifications of the environment. Many times, the decision between utilizing cloud native versus third-party solutions comes down to the actual requirements of the use case and the availability of the solution that fits those requirements.

Review Questions

1. Resource policies
2. Secrets management
3. Buckets
4. Public subnet
5. CASB
6. SWG
7. Application
8. Cloud native

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Chapter 23

Do I Know This Already?

1. B. An identity provider (IdP) is a service provider that also manages the authentication and authorization process on behalf of other systems in the federation.

2. B. Authentication is the process of proving the identity of a subject or user. Once a subject has identified itself in the identification step, the enforcer has to validate the identity—that is, be sure that the subject (or user) is the one it is claiming to be. This is done by requesting that the subject (or user) provide something that is unique to the requestor. This could be something known only by the user, usually referred to as *authentication by knowledge*, or owned only by the user, usually referred to as *authentication by ownership*, or it could be something specific to the user, usually referred to as *authentication by characteristic*.

3. A. Most certificates are based on the X.509 standard, which is a common PKI standard developed by the ITU-T that often incorporates the single sign-on (SSO) authentication method.

4. D. SAML is an open standard for exchanging authentication and authorization data between identity providers. SAML is used in many single sign-on (SSO) implementations.

5. D. A service account is typically used on a server to provide a separate set of credentials and permissions to an application or service that is running. For instance, a server that is running Apache web server might have an apache_ user account. That account would then be provided with only the access it needs to be able to perform the functions that services provide.

6. B. A tool you can use to validate who is logged in to a Windows system is the PsLoggedOn application. For this application to work, it has to be downloaded and placed somewhere on your local computer that will be remotely checking hosts. After it's installed, simply open a command prompt and execute the command **C:\PsTools\psloggedon.exe\\HOST_TO_CONNECT**.

7. C. For Linux machines, various commands can show who is logged in to a system, such as the **w**, **who**, **users**, **whoami**, and **last "user name"** commands.

8. B. Geolocation is the actual method of determining the physical location of the user trying to authenticate.

9. C. A time-based attribute can be used when authenticating and authorizing a user. When a user logs in and provides his or her identity, that user can be given specific access based on the time he or she connects. For instance, if you

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

do not expect that someone should be connecting to your wireless network at 3 a.m., then you can set a policy that blocks access between specific hours of operation.

10. B. Geotagging is the process of attaching location information in the metadata of files, such as pictures taken from a smartphone.

Review Questions

1. **whoami**
2. Biometrics
3. Two
4. One-time password (OTP)
5. Attribute-based access control (ABAC)
6. Geolocation

Chapter 24

Do I Know This Already?

1. A. Password keys are a technology typically deployed by corporations when implementing two-factor authentication. The primary use case is remote access to the organization's environment. However, many organizations also use them internally. These keys are especially important to use for access to highly sensitive data or applications that serve that data—for instance, financial applications or anything involving intellectual property, such as source code. Many types of password keys are on the market these days. They come in various form factors and are utilized in different ways. For instance, some are used by inserting into a computer USB port. Some are simply one-time password tokens used as a second factor when authenticating, whereas others are a combination of different functions.
2. B. A password vault is also often referred to as a password manager. It is simply a piece of software that is utilized to store and manage credentials. Typically, the credentials are stored in an encrypted database. Having an encrypted database protects the credentials from being compromised if the database is obtained by a threat actor through the compromise of a system holding the database.
3. A. NIST defines knowledge-based authentication (KBA) as authentication of an individual based on knowledge of information associated with his or her claimed identity in public databases. Knowledge of such information is considered to be private rather than secret, because it may be used in contexts

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

other than authentication to a verifier, thereby reducing the overall assurance associated with the authentication process. A popular use case for this type of authentication is to recover a username or reset a password. Typically, a set of predetermined questions is asked of the user. These questions must have already been provided by the end user at the time of account setup or provided as an authenticated user at a later time. The idea is that the information that was provided is something that only the user would know. That is why it is important to utilize a set of questions that cannot be easily guessed or are not public knowledge.

4. D. Using single sign-on (SSO), a user can log in once but gain access to multiple systems without being asked to log in again. This system is complemented by single sign-off, which is basically the reverse; logging off signs off a person from multiple systems. SSO is meant to reduce password fatigue or password chaos, which occurs when a person can become confused and possibly even disoriented when having to log in with several different usernames and passwords. This system is also meant to reduce IT help desk calls and password resets.

5. D. Single sign-on (SSO) is a derivative of federated identity management (also called FIM or FIdM). In this system, a user's identity and attributes are shared across multiple identity management systems. These various systems can be owned by one organization; for example, Microsoft offers the Forefront Identity Manager software, which can control user accounts across local and cloud environments.

6. B. Security Assertion Markup Language (SAML) is an open standard for exchanging authentication and authorization data between identity providers. SAML is used in many single sign-on (SSO) implementations.

   The OASIS Security Assertion Markup Language standard is currently the most-used standard for implementing federated identity processes. SAML is an XML-based framework that describes the use and exchange of SAML assertions in a secure way between business entities. The standard describes the syntax and rules to request, create, use, and exchange these assertions.

7. B. Discretionary access control (DAC) is an access control policy generally determined by the owner. Objects such as files and printers can be created and accessed by the owner. Also, the owner decides which users are allowed to have access to the objects and what level of access they may have. The levels of access, or permissions, are stored in access control lists (ACLs).

8. B. Mandatory access control (MAC) is an access control policy determined by a computer system, not by a user or owner, as it is in DAC. Permissions are predefined in the MAC model. Historically, it has been used in highly classified government and multilevel military systems, but you can find lesser

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

implementations of it in today's more common operating systems as well. The MAC model defines sensitivity labels that are assigned to *subjects* (users) and *objects* (files, folders, hardware devices, network connections, and so on).

9. B. Least privilege is a security principle in which users are given only the privileges needed to do their job and not one iota more. A basic example would be the Guest account in a Windows computer. This account (when enabled) can surf the web and use other basic applications but cannot make any modifications to the computer system. However, least privilege as a principle goes much further. One of the ideas behind the principle is to run the user session with only the processes necessary, thus reducing the amount of CPU power needed.
10. B. Role-based access control (RBAC) is an access model that, like MAC, is controlled by the system, and, unlike DAC, not by the owner of a resource. However, RBAC is different from MAC in the way that permissions are configured. RBAC works with sets of permissions instead of individual permissions that are label-based. A set of permissions constitutes a role. When users are assigned to roles, they can then gain access to resources.

Review Questions

1. Least privilege
2. Implicit deny
3. Linux
4. Role-based access control (RBAC)
5. Discretionary access control (DAC)
6. Privileged access management (PAM)
7. Mandatory access control (MAC)
8. Kerberos
9. Attribute-based access control (ABAC)
10. Hardware security module (HSM)
11. Knowledge-based authentication (KBA)

Chapter 25

Do I Know This Already?

1. A. Users need a private key to encrypt the digital signature of a private email. The difference in type of key is the level of confidentiality. A public key certificate obtained by a web browser is public and might be obtained by thousands

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

of individuals. The private key used to encrypt the email is not to be shared with anyone.

2. B. A public key certificate obtained by a web browser is public and might be obtained by thousands of individuals. A private key used to encrypt email is not to be shared with anyone.

3. A. Most certificates are based on the X.509 standard, which is a common PKI standard developed by the ITU-T that often incorporates the single sign-on (SSO) authentication method.

4. D. Components of an X.509 certificate include the following:

   - Owner (user) information, including public key
   - Certificate authority information, including name, digital signature, serial number, issue and expiration dates, and version

5. D. Many companies have subdomains for their websites. Generally, if you connect to a secure website that uses subdomains, a single certificate allows for connections to the main website and the subdomains. This is known as a wildcard certificate; for example, *.h4cker.org, meaning all subdomains of h4cker.org.

6. B. By modifying the Subject Alternative Name (SAN) field, an organization can specify additional hostnames, domain names, IP addresses, and so on.

7. B. Canonical Encoding Rules (CER) is a restricted version of BER in that it allows the use of only one encoding type; all others are restricted.

8. B. If a root CA is compromised, all of its certificates are then also compromised, which could affect an entire organization and beyond. The entire certificate chain of trust can be affected. One way to add a layer of security to avoid root CA compromise is to set up an offline root CA. Because it is offline, it cannot communicate over the network with the subordinate CAs, or any other computers for that matter. Certificates are transported to the subordinate CAs physically using USB flash drives or other removable media. Of course, you would need to have secure policies regarding the use and transport of media, and would need to incorporate data loss prevention (DLP), among other things. But the offline root CA has some obvious security advantages compared to an online root CA. Consider this offline mindset when dealing with critical data and encryption methods.

9. B. One way to add security to a certificate validation process is to use certificate pinning, also known as SSL pinning or public key pinning. It can help detect and block many types of on-path attacks by adding an extra step beyond normal X.509 certificate validation. Essentially, a client obtains a certificate from a CA in the normal way but also checks the public key in the server's certificate against a hashed public key used for the server name. This functionality

must be incorporated into the client side, so it is important to use a secure and up-to-date web browser on each client in order to take advantage of certificate pinning.

Review Questions

1. In X.509, the owner does not use a symmetric key.
2. A digital certificate includes the certificate authority's digital signature and the user's public key. A user's private key should be kept private and should not be within the digital certificate.
3. Decentralized. When creating key pairs, PKI has two methods: centralized and decentralized. In centralized, keys are generated at a central server and are transmitted to hosts. In decentralized, keys are generated and stored on a local computer system for use by that system.
4. Certificate revocation lists are digitally signed by the certificate authority for security purposes. If a certificate is compromised, it will be revoked and placed on the CRL. CRLs are later generated and published periodically.
5. The public key infrastructure is based on the asymmetric encryption concept.
6. You should implement a certificate revocation list so that stolen certificates, or otherwise revoked or held certificates, cannot be used.
7. A compromised certificate should be published to the certificate revocation list.
8. Public key encryption to authenticate users and private keys to encrypt the database. PKI uses public keys to authenticate users. If you are looking for a cryptographic process that allows for decreased latency, then symmetrical keys (private) would be the way to go. So, the PKI system uses public keys to authenticate the users, and the database uses private keys to encrypt the data.
9. A key escrow is implemented to secure a copy of the user's private key (not the public key) in case it is lost.
10. The browser must present the public key, which is matched against the CA's private key.

Chapter 26

Do I Know This Already?

1. A. Using the Linux **traceroute** command enables you to document hosts' locations on your local network and map out the current location/configuration and connected devices. A baseline network diagram should be used and continuously updated to document systems.

2. B. **nslookup** is a simple but practical command-line tool. It is principally used to find the IP address that corresponds to a host or the domain name that corresponds to an IP address (a process called reverse DNS lookup).
3. C. The Linux-centric **head** command reads the first 10 lines of a given filename.
4. B. The **cat** command is a widely used and universal tool. It copies standard input to standard output. The command supports scrolling if the text file doesn't fit the current screen.
5. B. A shell can have two types of variables: environment variables that are exported to all processes spawned by the shell and shell (local) variables.
6. D. The environment is an area that the shell builds every time it starts a session. This area contains variables that define system properties.
7. C. Wireshark is available on Windows, Linux, and macOS. There are both command-line and GUI versions.
8. D. The **pps-multi** command sets the number of packets to send for each time interval. This option must appear in combination with the following option: **pps**. This option takes an integer number as its argument. The value of the number is constrained to being greater than or equal to one. The default number for this option is one; therefore, you should use **--pps-multi=#** to set pps.
9. B. You can use the **dd** tool to create a complete image of the hard disk /dev/hda by using **# dd if = /dev/hda of = ~/hdadisk.img**. This image can be used to preserve forensic evidence of a computer system that was attacked or used to exploit other systems.
10. D. Three of the most notable exploitation frameworks are Metasploit, Core Impact, and Immunity Canvas, although there are a number of less famous frameworks.
11. C. Hashcat is one of the most popular and widely used password crackers in existence. It is available on every operating system and supports more than 300 different types of hashes. It enables highly parallelized password-cracking capabilities, enabling you to crack multiple different passwords on multiple different devices at the same time and to support a distributed hash-cracking system via overlays. Cracking is optimized with integrated performance tuning and temperature monitoring.
12. C. There are three methods to achieve data sanitization: physical destruction, cryptographic erasure, and data erasure. However, the downside of these techniques is that they damage the storage media and do not allow it to be sold or reused. The only real method to ensure the data is unrecoverable is to shred the hard drive or data as it is.

Review Questions

1. In computing, **traceroute** is a computer network diagnostic command for displaying possible routes and measuring transit delays of packets across a network.
2. Cuckoo Sandbox is open-source software for automating analysis of suspicious files. To do so, it makes use of custom components that monitor the behavior of the malicious processes while running in an isolated environment. You can throw any suspicious file at it, and in a matter of minutes, Cuckoo will provide a detailed report outlining the behavior of the file when executed inside a realistic but isolated environment.
3. **ping** verifies IP-level connectivity to another TCP/IP computer by sending Internet Control Message Protocol (ICMP) echo request messages.
4. **hping** supports TCP, UDP, ICMP, and RAW-IP protocols; has a traceroute mode; can send files between a covered channel, and provides many other features. It has a wide range of additional uses, including firewall testing, manual path MTU discovery, advanced traceroute, remote OS fingerprinting, advanced port scanning, remote uptime guessing, and TCP/IP stack auditing.
5. The **curl** command-line tool can transfer data to or from a server, using any of the supported protocols: HTTP, FTP, IMAP, POP3, SCP, SFTP, SMTP, TFTP, TELNET, LDAP, or FILE. **Curl** is powered by Libcurl.

Chapter 27

Do I Know This Already?

1. A. Playbook documents are step-by-step procedures and should be high level and focused on specific areas such as malware, insider threats, unauthorized access, ransomware, and phishing.
2. D. Data breach notification laws are becoming more common: the European Union's General Data Protection Regulation (GDPR), for instance, requires that companies report data security incidents within 72 hours of discovery.
3. C. A triage matrix provides an understanding of the severity of an incident so that it can be prioritized quickly and correctly.
4. C. The lessons learned phase of the incident response process should be performed no later than two weeks from the end of the incident. A Post Incident Response (PIR) meeting ensures information is fresh on the team's mind.
5. A. The tabletop exercise is often used to validate and/or improve an organization's incident response (IR) plan. Real-life scenarios are used to put the response plan to the test, highlighting areas where your team excels and areas to be addressed.

6. C and D. The Diamond Model places the basic components of malicious activity at one of the four points on a diamond shape: adversary, infrastructure, capability, and victim. The model provides for analysis of threats related to intrusion.

7. B. *Procedures* describe the way adversaries implement a technique. A procedure concerns the particular instance use and can be useful for understanding exactly how the technique is used and for replication of an incident with adversary emulation as well for specifics on how to detect the instance in use.

8. D. There are five key stakeholders for any IR team; they are IT Services, Security Management, Legal, Human Resources, and Public Relations.

9. C. When your team is engaged with an incident, you should have them set up proactive alerting. They don't need to call everyone every time, but your handlers need to plan ahead. Your incident response team needs to keep key contacts up to date so that when they have to notify contacts, it doesn't come as a surprise. Notifying key contacts with only incident-relevant data as it becomes available reduces overcommunication of nonimportant data.

10. C. A disaster recovery (DR) plan is a formal document created by organizations that contains detailed instructions on how to respond to unplanned incidents such as natural disasters, power outages, cyber attacks, or other disruptive events.

11. C. Disruptions lead to lost revenue, brand damage, and dissatisfied customers. The longer the recovery time, the greater the adverse impact to the business.

12. C. Many people think a disaster recovery (DR) plan is the same as a business continuity plan (BCP), but a DR plan focuses mainly on restoring an IT infrastructure and operations after a crisis. It's actually just one part of a complete business continuity plan, because a BCP looks at the continuity of the entire organization.

13. A. A continuity of operations plan (COOP) ensures the restoration of organizational functions in the shortest possible time, even if services resume at a reduced level of effectiveness or availability.

14. A, B, D. An incident response team is a group of people who prepare for and respond to any emergency incident, such as a natural disaster or an interruption of business operations.

15. B. NIST SP 800-53 outlines the requirements that contractors and federal agencies need to take to meet the Federal Information Security Management Act (FISMA). It requires data retention for a minimum of three years.

Review Questions

1. You should regularly test and update your incident response plan. Everyone who is part of the plan should understand their role and the role of others to help reduce confusion during a real event.
2. Regardless of how you choose to eradicate an infection, you need to have a plan for increased monitoring of any affected systems for some period of time after the eradication process within 30 days.
3. Incident response simulations are internal events that provide a structured opportunity to practice your incident response plan and procedures during a realistic scenario. SIRS events are fundamentally about being prepared and iteratively improving your response capabilities.
4. The Diamond Model of Intrusion Analysis emphasizes the relationships and characteristics of four basic components: the adversary, capabilities, infrastructure, and victims.
5. Privilege escalation: Attackers often need more privileges on a system to get access to more data and permissions: for this, they need to escalate their privileges, often to an Admin.

Chapter 28

Do I Know This Already?

1. A. Historical vulnerability scans can provide significant insight after an incident. By comparing the previous scans with the most recent, you can also look for variances in devices and systems that may have been changed.
2. D. Network vulnerability scans should include all devices with an IP address (workstations, laptops, printers and multifunction printers, IoT devices, routers, switches, hubs, IDS/IPS, servers, wireless networks, and firewalls) and all the software running on them.
3. B. Sensors should be deployed before an incident. Sensor placement around the network allows for greater visibility and can aid in forensic investigation by quickly identifying the depths of the spread.
4. C. SIEMs can tune sensitivity to what is considered suspect behavior or suspicious files (risk-based prioritization) to help reduce or increase the amount of data/matches during an investigation.
5. D. The log data you collect from your systems and devices may seem pretty routine. These logs could contain the precise evidence needed to investigate and successfully eradicate incidents from your network.

6. D. Session Initiation Protocol (SIP) is a signaling protocol used to establish, maintain, and tear down a call when terminated. SIP allows the calling parties' called user agents to locate one another using the network.

7. A, B, C. Syslog is available on most network devices (for example, routers, switches, and firewalls), as well as printers and Linux-based systems. A syslog server listens for and then logs data messages coming from syslog clients. Rsyslog and syslog-ng build on syslog capabilities by adding support for advanced filtering and configuration. Event Viewer is a Windows-based tool that enables users or administrators to view event logs on Windows-based remote or local systems.

8. C. If you need to access logs for a specific time window—for instance, **journalctl -since 2021-03-15 15:05:00**—the time format is YYYY-MM-DD HH:MM:SS. **journalctl** is a valuable tool used to collect logs and sort through mountains of data to help you find a needle in a haystack.

9. A. NXLog can process high volumes of event logs from many different sources. Log processing includes rewriting, correlating, alerting, filtering, pattern matching, log file rotation buffering, and prioritized processing. Application, system, and security logs are Windows-specific Event Viewer logs.

10. C. Bandwidth monitors track bandwidth use over all areas of the network, including devices, applications, servers, WAN, and Internet links. One benefit of deploying bandwidth monitors is that they map out historical trends for capacity planning. With bandwidth monitors, you can quickly identify abnormal bandwidth usage, top talkers, and unique communications, all useful in finding infected systems that may be exfiltrating data or scanning the network looking to spread to other hosts.

11. A. Metadata is created from every activity you perform, whether it's on your personal computer or online, every email, web search, and social or public application. Metadata is defined as "data that provides information about other data." On its own, it might not seem like much, but when combined with additional context can lead to a break in a case.

12. C. A flow is a unidirectional set of packets sharing common session attributes, such as source and destination, IP, TCP/UDP ports, and type of service. NetFlow statefully tracks flows or sessions, aggregating packets associated with each flow into flow records, which are bidirectional flows.

13. D. Protocol analyzers allow network engineers and security teams to capture network traffic and perform analysis of the captured data to identify potential malicious activity or problems with network traffic. The network traffic data can be observed in real time for troubleshooting purposes, monitored by an

alerting tool such as a SIEM to identify active network threats, and/or retained to perform forensic analysis.

Review Questions

1. Indicators can be anything from additional TCP/UDP ports being shown as open to detection of unauthorized software, or scheduled host system events and even unrecognized outbound communications.
2. Data correlation allows you to take data and logs from disparate systems, like Windows server events, firewalls logs, VPN connections, and RAS devices, and bring them all together to see exactly what took place during that event.
3. Logging for critical process information about user, system, and web application behavior can help incident responders build a better understanding of what normal looks like when an application is running and being used.
4. The DNS protocol has two message types: queries and replies. Both use the same format. These messages are used to transfer resource records (RRs). An RR contains a name, a time-to-live (TTL), a class (normally IN), a type, and a value. There are nearly a dozen different types of logs that are of particular interest; obtaining and including them in your investigation can help build a full picture.
5. VoIP technology is an attractive platform to criminals. The reason is that call managers and VoIP systems are global telephony services, in which it is difficult to verify the user's location and identification.

Chapter 29

Do I Know This Already?

1. A. You spend time up front building a whitelist of approved applications for the application approved list. Then with your central management or endpoint security solution, you roll out the whitelist enterprisewide to all endpoints.
2. B. An application block list or deny list is a basic access control mechanism that denies specific applications or code on the list from being installed or run.
3. D. The main action of the quarantine function is to safely store reported objects such as malware, infected files, or potentially unwanted applications. During an investigation, this storage log should be one of the places investigators check for evidence and history of suspect files.
4. D. With mobile device management (MDM), it is important to establish a baseline understanding of the changes that are common to specific types of

mobile devices. Mobile platforms have several attack vectors that you need to consider: hardware, firmware, mobile OS, applications, and the device combination.

5. D. Data loss prevention (DLP) is an end-to-end goal that ensures users do not send sensitive or critical information outside the corporate network. The term routinely describes software products that help a network administrator control the data that users can view or transfer. Intellectual property, corporate data, and customer data are some of the types of data you would use DLP to help protect against exfiltration.
6. A. Once a compromised endpoint has been infected because malware, a virus, or a Trojan was detected, or it is part of a much wider attack, it should be quarantined and isolated.
7. C. When you know where the issue is and what systems have been affected, you should contain those systems so that threats cannot spread through the network. To do this, you disable network access for the affected computers and devices, or you place them in a sandbox network. You also should change passwords so intruders no longer have access.
8. D. Network segmentation is an architectural approach that divides a network into multiple networks, subnets, or segments. It allows network administrators to control the flow of traffic between these networks based on a granular policy.
9. A. A SOAR system alerts for suspected phishing emails, endpoint attacks, failed user logins, malware, and other threat information that come from a variety of detection sources, such as SIEMs, systems, switches, and logging services. IT operations use a SOAR runbook for reference for routine procedures that administrators perform. A SOAR playbook provides manual orchestration of incident response. Web content filtering is the practice of blocking access to web content that may be deemed offensive or inappropriate, or even contain dangerous items.

Review Questions

1. The purpose is to specify an index of approved software applications or executable files that are permitted to be present and active on a computer system.
2. When a file is quarantined, the file is moved to a location on disk where it cannot be executed.
3. A set of tools and processes is used to ensure that sensitive data is not lost, misused, or accessed by unauthorized users. These tools allow only authorized persons to have access and to run copy/move commands on those specific files.

4. Certificate revocation is the act of invalidating a certificate before its scheduled expiration date. A certificate should be revoked immediately when its private key shows signs of being compromised.
5. A runbook consists of a series of conditional steps to perform actions such as enriching data, containing threats, and sending notifications automatically as part of the incident response or security operations process.

Chapter 30

Do I Know This Already?

1. D. Data preservation is the first step when litigation has been filed or will be soon filed, with a focus on the preservation of data in its current state, such as emails, SMS, MMS, and deleted messages (still on disk and not destroyed) on all devices including cell phones, PCs, and mobile devices.
2. B. File formats can vary, as well as settings for recording video, such as frames per second (FPS) and video resolution. These features can all factor into how and what video information is stored.
3. B. It has become imperative that evidence collection standards and procedures are consistent, documented, and coherent in the chain of custody. Strict guidelines are to be followed, and accurate documentation must be kept.
4. B. Evidence tagging helps identify collected items. A tag can consist of something as little as a sticker with the date, time, control number, and name or initials of the investigator. Using a control number is an easy way to identify a piece of evidence in documentation, such as a chain of custody. The combination of a tag and photographs can provide the exact location and condition in which the item was collected.
5. A. The collection of evidence should start with the most volatile item and end with the least volatile. The order of volatility is the order in which the digital evidence is collected.
6. D. Forensic artifacts are objects that have forensic value. They are not behaviorally driven; they do not necessarily reflect the behavior or intent of a threat actor or adversary. Some of the artifacts that can be extracted from suspect hosts are logs, the registry, the browser history, the RDP cache, and Windows Error Reporting (WER).
7. B. RAM is considered volatile memory. It is perceived to be more trusted than nonvolatile memory, like ROM, disk, magnetic, or optical storage. Investigations using live forensic techniques require special handling because the volatile data in RAM can contain code used by attackers.

8. D. Checksums may also be called hashes. Small changes in a file produce different-looking checksums. You can use checksums to check files and other data for errors that occur during transmission or storage, as well as for evidence in a forensic investigation to ensure it hasn't been tampered with.

9. B. *Cloud-based* and *on-premises* are simply terms that describe where systems store data. Many of the same vulnerabilities that affect on-premises systems also affect cloud-based systems.

10. D. A copy of digital evidence must be properly preserved and collected in accordance with forensic best practices. Otherwise, the digital evidence may be inadmissible in court, or spoliation sanctions may be imposed.

11. A. Organizations should have a legal hold process to perform e-discovery to preserve and gather such information.

12. C. Forensic data recovery is the extraction of data from damaged, deleted, or purposely destroyed evidence sources in a forensically sound manner. This method of recovering data means that any evidence resulting from it can later be relied on in a court of law.

13. D. Nonrepudiation makes it difficult to successfully deny who and where a message came from as well as the authenticity and integrity of that message. Digital signatures can offer nonrepudiation when it comes to online transactions.

14. A. Counterintelligence is information gathered and activities conducted to protect against espionage, other intelligence activities, or sabotage conducted by or on behalf of other elements. The intelligence is designed to quickly direct resources to the most significant problems first and address them head on.

Review Questions

1. Whether evidence is admissible is determined by following three rules: (1) Best evidence means that courts prefer original evidence rather than copies to avoid alteration of evidence. (2) The exclusionary rule means that data collected in violation of the Fourth Amendment (no unreasonable searches or seizures) is not admissible. (3) Hearsay is second-hand evidence and is often not admissible, but some exceptions apply.

2. It must be (1) sufficient, which is to say convincing without question; (2) competent, which means it is legally qualified; and (3) relevant, which means it must matter to the case at hand.

3. Computers use checksum-style techniques to check data for problems in the background. You could also use checksums to verify the integrity of any other type of file, from applications to documents and media. Forensic investigators

use checksums to ensure data is not tampered with after it has been collected from an incident.

4. By definition, forensic copies are exact, bit-for-bit duplicates of the original. To verify this, you can use a hash function to produce a type of unique checksum of the source data. Hash functions have four defining properties that make them useful; they are deterministic, collision resistant, pre-image resistant, and computationally efficient.

5. Network forensic analysis tools (NFATs) typically provide the same functionality as packet sniffers, protocol analyzers, and SIEM software in a single product. NFAT software focuses primarily on collecting, examining, and analyzing network traffic.

Chapter 31

Do I Know This Already?

1. A. Managerial controls are techniques and concerns addressed by an organization's management (managers and executives). Generally, these controls focus on decisions and the management of risk. They also concentrate on procedures, legal and regulatory policies, the software development lifecycle (SDLC), the computer security lifecycle, information assurance, and vulnerability management/scanning. In short, these controls focus on how the security of your data and systems is managed.

2. B. Operational controls are the controls executed by people. They are designed to increase individual and group system security. They include user awareness and training, fault tolerance and disaster recovery plans, incident handling, computer support, baseline configuration development, and environmental security. The people who carry out the specific requirements of these controls must have technical expertise and understand how to implement what management desires of them.

3. A. Technical controls are the logical controls executed by the computer system. Technical controls include authentication, access control, auditing, and cryptography. The configuration and workings of firewalls, session locks, RADIUS servers, or RAID 5 arrays would be within this category, as well as concepts such as least privilege implementation.

4. D. Preventative controls are employed before an event and are designed to prevent an incident. Examples include biometric systems designed to keep unauthorized persons out, network intrusion prevention systems (NIPSs) to prevent malicious activity, and RAID 1 to prevent loss of data. They are also sometimes referred to as deterrent controls. Preventative controls enforce

security policy and should prevent incidents from happening. The only way to bypass a preventative control is to find a flaw in its implementation or logic. These controls are usually not optional. Examples of preventative controls are access lists, passwords, and fences.

5. D. Detective controls aim at monitoring and detecting any unauthorized behavior or hazard. These types of controls are generally used to alert of a failure in other types of controls such as preventative, deterrent, and compensating controls. Detective controls are very powerful while an attack is taking place, and they are useful in the post-mortem analysis to understand what has happened. Audit logs, intrusion detection systems, motion detection, and Security Information and Event Management (SIEM) systems are examples of detective controls.

6. B. Corrective controls are used after an event. They limit the extent of damage and help the company recover from damage quickly. Tape backup, hot sites, and other fault tolerance and disaster recovery methods are also included here. They are sometimes referred to as compensating controls. Corrective controls include all the controls used during an incident to correct the problem. Quarantining an infected computer, sending a guard to block an intruder, and terminating an employee for not having followed the security policy are all examples of corrective controls.

7. B. Compensating controls, also known as alternative controls, are mechanisms put in place to satisfy security requirements that are either impractical or too difficult to implement. For example, instead of using expensive hardware-based encryption modules, an organization might opt to use network access control (NAC), data loss prevention (DLP), and other security methods. Or, on the personnel side, instead of implementing separation of duties, an organization might opt to do additional logging and auditing. You should approach compensating controls with great caution. They do not give the same level of security as their replaced counterparts.

Review Questions

1. Managerial
2. Operational
3. Technical
4. Preventative
5. Deterrent
6. Detective

7. Corrective
8. Compensating
9. Physical
10. Physical

Chapter 32

Do I Know This Already?

1. A. The General Data Protection Regulation is a European Union (EU) law that was enacted in 2018. Its overall focus is on data protection and privacy for individuals. Although it is a law enacted in the EU, it applies to any organization collecting information about people in the EU. This means that if your organization collects and handles the personal data of EU citizens, then this regulation applies to you. For instance, if you run a business that offers goods or services in the EU and that business requires you to collect information about your customers, then you would be required to follow the GDPR requirements. Not following them could result in large fines. This is one of the factors that makes GDPR a larger concern to organizations than many other laws that have been in place for many years. GDPR penalties can be very high. For additional information on GDPR, refer to https://gdpr.eu/.
2. B. The Sarbanes–Oxley Act (SOX), enacted in 2002, governs the disclosure of financial and accounting information.
3. A. The Health Insurance Portability and Accountability Act (HIPAA), enacted in 1996, governs the disclosure and protection of health information.
4. D. The Center for Internet Security (CIS) is a nonprofit organization that was established in 2000. Its overall goal is to provide security best practice guidance for enhancing the security of cyberspace.
5. C. The National Institute of Standards and Technology (NIST) developed the Risk Management Framework (RMF) in 2017 as a result of an executive order from the president, which required all federal agencies to comply with it. For mor information about the NIST RMF, visit www.nist.gov/cyberframework/risk-management-framework.
6. C. The National Institute of Standards and Technology (NIST) developed the Cybersecurity Framework (CSF) in 2014 as a result of executive order 13636 from the president. The NIST CSF is made of five core functions: Identify, Protect, Detect, Respond, and Recover. For more information about the NIST CSF, visit www.nist.gov/cyberframework.

7. A. The Cloud Security Alliance (CSA) is a nonprofit organization established in 2008 with the goal of promoting security best practices in cloud computing environments. The Cloud Controls Matrix is a framework established by the CSA for cloud computing. This organization also developed the reference architecture to help cloud providers with guidance on developing secure interoperability best practices. For more information related to CSA, visit https://cloudsecurityalliance.org/.
8. C. Security Content Automation Protocol (SCAP) was created to provide a standardized solution for security automation. The SCAP mission is to maintain system security by ensuring security configuration best practices are implemented in the enterprise network, verifying the presence of patches, and maintaining complete visibility of the security posture of systems and the organization at all times.

Review Questions

1. General Data Protection Regulation (GDPR)
2. Sarbanes-Oxley Act (SOX)
3. Health Insurance Portability and Accountability Act (HIPAA)
4. Center for Internet Security
5. National Institute of Standards and Technology
6. Cloud Security Alliance
7. Security Content Automation Protocol (SCAP)

Chapter 33

Do I Know This Already?

1. A. The Privacy Act of 1974 sets many standards when it comes to the security of personally identifiable information (PII). However, most organizations will go further and define their own privacy policy, which explains how users' identities and other similar information will be secured. For example, if an organization has an Internet-based application that internal and external users access, the application will probably retain some of their information—possibly details of their identity. Not only should this information be secured, but the privacy policy should state in clear terms what data is allowed to be accessed, and by whom, as well as how the data will be retained and distributed (if at all). An organization might also enact a policy that governs the labeling of data to ensure that all employees understand what data they are handling and to prevent the mishandling of confidential information. Before any system

administrators or other personnel gather information about these users, they should consult the privacy policy.

2. B. Acceptable use policies (AUPs) define the rules that restrict how a computer, network, or other system may be used. They state what users are and are not allowed to do when it comes to the technology infrastructure of an organization. Often, employees must sign an AUP before they begin working on any systems. This policy protects the organization but also defines to employees exactly what they should and should not be working on.

3. A. Separation of duties defines when more than one person is required to complete a particular task or operation. This distributes control over a system, infrastructure, or particular task.

4. D. When it comes to information security, due diligence means ensuring that IT infrastructure risks are known and managed. An organization needs to spend time assessing risk and vulnerabilities and might state in a policy how it will give due diligence to certain areas of its infrastructure.

5. D. Due care is the mitigation action that an organization takes to defend against the risks that have been uncovered during due diligence.

6. B. Due process is the principle that an organization must respect and safeguard personnel's rights. The purpose is to protect the employee from the state and from frivolous lawsuits.

7. A. All employees should be trained on personally identifiable information (PII). This information is used to uniquely identify, contact, or locate a person. This type of information could be a name, birthday, Social Security number, biometric information, and so on. Employees should know what identifies them to the organization and how to keep that information secret and safe from outsiders. Another key element of user education is the dissemination of the password policy. Employees should understand that passwords need to be complex and know the complexity requirements. They should also understand that they should never give out their password or ask for another person's password to any resource.

8. C. A memorandum of understanding (MOU) is not an agreement at all, but an understanding between two organizations or government agencies. It does not specify any security controls either. However, a memorandum of *agreement* (MOA) does constitute a legal agreement between two parties wishing to work together on a project but still does not detail any security controls.

9. B. A business partnership agreement (BPA) is a type of contract that can establish the profits each partner will get, what responsibilities each partner will

have, and exit strategies for partners. Often this type of agreement applies to supply chain and business partners.

10. B. Top secret means the highest sensitivity of data; few people should have access, and security clearance may be necessary. Information is broken into sections on a need-to-know basis.

11. B. When you're discussing policies for systems internal to your organization or devices that are owned by your organization, such as servers and laptops, it is also important to detail policies regarding how they should be deployed. For instance, service accounts should follow a specific credential policy. Often service accounts have higher-level permissions to enable the specific service. That means these accounts can cause more damage if they are compromised. They should always be configured using the least privilege approach. There should also be a service account for each service.

12. B. Change management is a structured way of changing the state of a computer system, network, or IT procedure. The idea behind this is that change is necessary, but that an organization should adapt with change and be knowledgeable of it. Any change that a person wants to make must be introduced to the heads of each department that it might affect. They must approve the change before it goes into effect. Before this happens, department managers will most likely make recommendations and/or give stipulations. When the necessary people have signed off on the change, it should be tested and then implemented. During implementation, it should be monitored and documented carefully.

Review Questions

1. Acceptable use policy (AUP)
2. Separation of duties
3. Code of ethics
4. Acceptable use policy (AUP)
5. Change management
6. Service-level agreement (SLA)
7. Public information
8. Job rotation
9. Guidelines and enforcement

Chapter 34

Do I Know This Already?

1. A. Risks to your organization or environment can come in many shapes and forms. The primary concern of most organizations is external risk. This is, of course, the biggest concern to most because it is risk that comes from an external entity who could have many different motivations and therefore targets the inside of your organization. Many imagine an external "hacker" as someone sitting in a basement hammering away at the keyboard in front of 10 different monitors in a black hoodie. That, of course, is usually not the case. External risk most likely is from an organized threat actor or an organization of threat actors who have various motivations. To carry out their attacks, they will use many different methods. The primary goal of external attackers is to gain access to your organization's computing environment, gain a foothold, and keep it as long as possible to carry out their objectives, whatever they may be.

2. B. Many organizations tend to overlook internal risks. The majority of internal risk stems from employees or those internal to the organization such as contractors. While the motivation of external cybercriminals may be to gain and keep access, the internal threat actor already has access to the organization's environment. This person's motivations are usually different from those of external threat actors—although the goal may be the same in the end. Most internal attacks result in the exfiltration or destruction of sensitive data that belongs to the organization. The theft of intellectual property (IP) is a primary goal of internal and external threat actors.

3. A. Risk management can be defined as the identification, assessment, and prioritization of risks, and the mitigating and monitoring of those risks. Specifically, when we talk about computer hardware and software, risk management is also known as information assurance (IA). The two common models of IA include the well-known CIA, and the DoD's "Five Pillars of IA," which comprise the concepts of the CIA triad (confidentiality, integrity, and availability) but also include authentication and nonrepudiation.

4. D. Some organizations opt to avoid risk. Risk avoidance usually entails not carrying out a proposed plan because the risk factor is too great. An example of risk avoidance: If a high-profile organization decided not to implement a new and controversial website based on its belief that too many attackers would attempt to exploit it.

5. C. An example of risk transference (also known as risk sharing) would be an organization that purchases cybersecurity insurance for a group of servers in a data center. The organization still takes on the risk of losing data in the case of server failure, theft, and disaster, but transfers the risk of losing the money those servers are worth in case they are lost.
6. B. Risk assessment is the attempt to determine the number of threats or hazards that could possibly occur in a given amount of time to your computers and networks.
7. B. Qualitative risk assessment is an assessment that assigns numeric values to the probability of a risk and the impact it can have on the system or network.
8. C. NIST defines risk mitigation as "prioritizing, evaluating, and implementing the appropriate risk-reducing controls/countermeasures recommended from the risk management process."
9. A. Person-made disasters can be defined as being caused by the influence of humans.
10. A. Mean time between failures (MTBF) defines the average number of failures per million hours of operation for a product in question.
11. A. Although it's impossible to predict the future accurately, it can be quantified on an average basis using concepts such as mean time between failures (MTBF).

Review Questions

1. Recovery plan
2. Recovery time objective (RTO)
3. Impact determination
4. Residual risk
5. Single point of failure
6. Mean time between failures (MTBF)
7. Quantitative risk assessment
8. Risk mitigation
9. Qualitative risk assessment
10. Risk assessment

Chapter 35

Do I Know This Already?

1. A. Regardless of what industry you are responsible for protecting, the ultimate goal is to protect your intellectual property. This is, of course, the crown jewels and can be in many different forms. For a company that develops software, the crown jewels are the source code of the product it is selling. Even a company that produces food products has intellectual property that it is trying to protect, such as a secret recipe.
2. B. The US Securities and Exchange Commission (SEC) requires that any publicly traded company provide a public notification and disclosure of a data breach.
3. A. Unauthorized access to top secret information would cause grave damage to national security.
4. D. Unauthorized access to unclassified information would cause no damage to national security.
5. D. Unauthorized access to private information could cause severe damage to the organization. Examples of information or assets that could receive this type of classification are human resources (HR) information (for example, employee salaries) and medical records.
6. B. According to the Executive Office of the President, Office of Management and Budget (OMB) and the U.S. Department of Commerce, Office of the Chief Information Officer, personally identifiable information, or PII, refers to "information which can be used to distinguish or trace an individual's identity."
7. B. *Data minimization* is a term used to explain a concept or approach to privacy design. The overall concept of data minimization is simply to minimize the amount of your personal information that is consumed by online entities. Data minimization is a privacy tool that is used in many different ways. For instance, a website may choose to not store your personal information if it is not needed—as opposed to many that store it and even resell it for a profit. Additionally, it can be used to develop policies regarding the amount of time the data that is collected about you is actually maintained before being permanently deleted. Individuals can also use tools that will clear information from applications such as web browsers. As you know, web browsers collect a large amount of data, which could in turn be compromised. Minimizing this data reduces the risk of such a compromise.

8. C. Another privacy enhancing technology concept is data masking. The goal of data masking is to protect or obfuscate sensitive data. This goal must be achieved while not rendering the data unusable in any way. An example of data masking being used in a real-world environment would be data that is displayed on terminal screens in banks or doctors' offices. Social Security numbers can be masked to show only the last four digits so that they can be used for verification purposes while not exposing the full Social Security number.

9. B. The data owner, also called the information owner, is usually part of the management team and maintains ownership of and responsibility over a specific piece or subset of data. Part of the responsibility of this role is to determine the appropriate classification of the information, ensure that the information is protected with controls, to periodically review classification and access rights, and to understand the risk associated with the information.

10. D. The General Data Protection Regulation (GDPR) in the European Union (EU) defines the information lifecycle in four different phases. It is sometimes named differently based on the source; however, the phases themselves are still the same. Starting with the collection of data, this is obviously the phase where the data is consumed by the data processor. GDPR states that when collecting data there must be a defined consent from the data owner as well as a clear definition of how the data will be used. The overall intent is to follow the principle of collecting only data that is necessary and not overcollecting.

11. A. Before data is collected, stored, secured, and disposed of throughout the information lifecycle, it is important to understand how that data, if compromised, could impact the privacy of the individuals whose data it holds. To accomplish this, an organization should complete an impact assessment on any new projects that are to be instated where data will be collected as well as any time the scope of the data use will change. These impact assessments are also sometimes called Privacy Impact Assessments (PIAs) or Data Privacy Impact Assessments (DPIAs). The result of an impact assessment should produce some kind of report that will identify specific high risks to the data subjects and provide recommendations on how that risk can be minimized. A PIA or DPIA is also something that is required in the GDPR.

12. D. Another privacy concept that has been adopted by the General Data Protection Regulation (GDPR) in the European Union (EU) is the terms of agreement. In many cases, this is called the data processing agreement. The overall purpose of the data processing agreement is to protect the personal information and the individuals the data is about. The agreement is actually a legal contract that is agreed upon by any entities that will fit the role of data processor in the information lifecycle. This is one of the basic requirements of GDPR and will lead to fines if not followed by an organization collecting data.

13. D. Along with the agreement of how data will be collected, utilized, and processed by an organization, the organization must also provide notification to the individuals it is collecting data from or about. Again, this is a requirement for the General Data Protection Regulation (GDPR) in the European Union (EU). GDPR ensures that individuals are notified about how their data is being used. This is done via a privacy notice. The notice itself is a document sent from the collecting organization stating how it is conforming to data privacy principles.

Review Questions

1. Private
2. Top secret
3. PII
4. PHI
5. Data minimization
6. Tokenization
7. Data controller
8. Data protection officer (DPO)

This page intentionally left blank

CompTIA Security+ (SY0-601) Cert Guide Exam Updates

Over time, reader feedback allows Pearson to gauge which topics give our readers the most problems when taking the exams. To assist readers with those topics, the authors create new materials clarifying and expanding on those troublesome exam topics. As mentioned in the Introduction, the additional content about the exam is contained in a PDF on this book's companion website, at http://www.pearsonitcertification.com/title/9780136770312.

This appendix is intended to provide you with updated information if CompTIA makes minor modifications to the exam upon which this book is based. When CompTIA releases an entirely new exam, the changes are usually too extensive to provide in a simple updated appendix. In those cases, you might need to consult the new edition of the book for the updated content. This appendix attempts to fill the void that occurs with any print book. In particular, this appendix does the following:

- Mentions technical items that might not have been mentioned elsewhere in the book
- Covers new topics if CompTIA adds new content to the exam over time
- Provides a way to get up-to-the-minute current information about content for the exam

Always Get the Latest at the Book's Product Page

You are reading the version of this appendix that was available when your book was printed. However, given that the main purpose of this appendix is to be a living, changing document, it is important that you look for the latest version online at the book's companion website. To do so, follow these steps:

Step 1. Browse to www. pearsonitcertification.com/title/9780136770312.

Step 2. Click the Updates tab.

Step 3. If there is a new Appendix B document on the page, download the latest Appendix B document.

NOTE The downloaded document has a version number. Comparing the version of the print Appendix B (Version 1.0) with the latest online version of this appendix, you should do the following:

- **Same version:** Ignore the PDF that you downloaded from the companion website.
- **Website has a later version:** Ignore this Appendix B in your book and read only the latest version that you downloaded from the companion website.

Technical Content

The current Version 1.0 of this appendix does not contain additional technical coverage.

Index

Symbols

../ (dot-dot-slash) attack 76, 274–275
; (semicolon) 73
' (single quotation mark) 73
_ (underscore) 740
0phtCrack 44
2FA (two-factor authentication) 298
5G communications 357–358
802.1X standard 510, 553–556, 562, 664–667, 673

A

A record (Address mapping record) 796
AAA (authentication, authorization, and accounting) framework 306
AAR (after action report) 928–929
ABAC (attribute-based access control) 638–645, 678, 679
acceptable use policies (AUPs) 898, 900
acceptance of risk 919
access control. *See also* 802.1X standard; identity; passwords
  access control entries (ACEs) 643
  access control lists (ACLs) 490, 528, 535, 643, 831
  attribute-based 638–645, 678, 679
  best practices 680–681
  centralized versus decentralized 679
  centralized/decentralized 640
  conditional access 678, 679
  delegation of access 662
  discretionary 674–676, 679
  identity and access management (IAM) 605, 633
  implicit deny 680
  least privilege 264, 630, 681, 908
  mandatory 676, 679
  network 510–511
  permissions 640–645
    cloud computing 605, 610
    inheritance 644–646
    open 150
    types of 646
  privileged access management (PAM) 678, 679
  role-based 677, 679
  rule-based 677, 678, 679
  summary of 679
  user access recertification 645
  vestibules 372–373
access control entries (ACEs) 643
access control lists (ACLs) 490, 528, 535, 643, 831
access points (APs)
  rogue 99
  security 562–563
accounting, AAA framework for 306
accounts 629–633. *See also* access control; passwords
  administrator 908
  auditing 635, 639
  harvesting 18
  permissions 640–645
    cloud computing 605, 610
    inheritance 644–646

open 150
types of 646
policies 633
root 908
service 908
ACEs (access control entries) 643
ACI (Application Centric Infrastructure) 243
acknowledgement (ACK) packets 84
ACLs (access control lists) 490, 528, 535, 643, 831
acquisition, forensic
artifacts 853
cache 852
checksums 857
data breach notification laws 855–856
definition of 847
device 850–851
disk 848
firmware 851
hashing 856–857
integrity 856
network 852–853
operating system 850
order of volatility 848
on-premises versus cloud 853–854
random-access memory (RAM) 848–849
regulatory and jurisdictional 855
right-to-audit clauses 854
snapshot 851–852
swap/pagefile 849–850
Active Directory (AD) 291–292
Active Directory Certificate Services (AD CS) utility 691
Active Directory Users and Computers (ADUC) 640
active reconnaissance 18, 204–205
active/active load balancing 488
active/passive load balancing 488
Activity Monitor 542
actors, threat
attack vectors 122–123
attributes of 122
types of 120–121
AD. *See* Active Directory (AD)
additional or associated data (AEAD) 404
Address mapping record (A record) 796
Address Resolution Protocol. *See* ARP (Address Resolution Protocol)
address space layout randomization (ASLR) 76, 265, 272
addresses
IPv4 443–444
IPv6 536–537
MAC (media access control) 101, 511
network address allocation 443–444
network address translation 501, 529, 562
virtual IP 488
administrator accounts 908
admissibility, evidence 843
ADSP (Author Domain Signing Practices) 110
ADUC (Active Directory Users and Computers) 640
Advanced Encryption Standard (AES). *See* AES (Advanced Encryption Standard)
Advanced IP scanner 721
advanced persistent threats (APTs) 35, 120–121, 451, 770
AE (authenticated encryption) 404
AEAD (additional or associated data) 404
aerospace application-embedded systems 348–350
AES (Advanced Encryption Standard) 412, 430, 475, 552
AES-GCM 498
AES-GMAC 498
AFL (American Fuzzy Lop) 269
after action report (AAR) 928–929
aggregation, log 186
aggregators 526

Agile development methodology 258–259
agreement, terms of 948
AH (Authentication Header) 437, 520
AI (artificial intelligence) 50–51
AI (Asset Identification) 885, 941–942
AICPA (American Institute of Certified Public Accountants) 883
AI/ML (artificial intelligence and machine learning) 50–51
AIR (As-if Infinitely Ranged) integer model 77
air gaps 384, 385
air traffic control (ATC) 349–350
aircraft systems 348–350
AirMagnet 99
AIS (automated indicator sharing) 125
aisles, hot/cold 386
alarms 374, 870
ALE (annualized loss expectancy) 922
alerts, SIEM 788
ALG (application-level gateway) 529
algorithms 50–51
  Grover's 402
  hashing 218–219, 856–857
    Digital Signature Algorithm (DSA) 396, 412
    Elliptic Curve Digital Signature Algorithm (ECDSA) 551–552
    Message Digest Algorithm 5 (MD5) 55, 219
    Secure Hash Algorithm (SHA) 55, 551–552
  key generation 395
  message authentication code (MAC) 410
  online resources 498
  public key 411
  scheduling 488
  Shor's 402
  signature verifying 395
  signing 395
allocation, network address 443–444
allow lists 467, 578, 583, 822
ALTER DATABASE statement 71
ALTER TABLE statement 71
Alureon rootkit 35–36
always-on VPN functionality 495
Amazon Web Services (AWS) 232–233, 244, 603, 853, 870
American Fuzzy Lop (AFL) 269
American Institute of Certified Public Accountants (AICPA) 883
amplification attacks 112
analytics logs 383
Android Auto 347
Angry IP scanner 721
annualized loss expectancy (ALE) 922
annualized rate of occurrence (ARO) 922
anomaly-based analysis 521–523
anonymization 945
anti-forensics 770
antimalware 452
antivirus software 451
anycast addresses 537
anything as a service (XaaS) 139, 232
AP isolation 562
Apache
  HTTP Server 146
  Mesos 240
  web servers 794
APIs (application programming interfaces) 86
  API-based keyloggers 42
  attacks 55, 85–86, 602
  definition of 240–241
  infrastructure as code 241–243
  inspection and integration 607, 610
  micro-segmentation 240–241
  security considerations 216
  Shodan 203–204
APP (Australia Privacy Principles) 220
Apple
  Apple Pay 462, 584
  CarPlay 347
  macOS Activity Monitor 542

appliances, network 513–514. *See also* firewalls
aggregators 526
hardware security modules (HSMs) 524
jump servers 514
network intrusion detection systems (NIDSs) 517–518
advantages/disadvantages 519–520
anomaly-based analysis 521–523
definition of 519–520
heuristic-based analysis 521
inline versus passive 523–524
promiscuous mode 517
signature-based 520–521
stateful pattern-matching recognition 521
network intrusion prevention systems (NIPSs)
advantages/disadvantages 519–520
anomaly-based analysis 521–523
definition of 518–520
false positives/false negatives 519
heuristic-based analysis 521
inline versus passive 523–524
signature-based 520–521
proxy servers 514–516
sensors 524–525
application allow lists. *See* allow lists
application block/deny lists. *See* block/deny lists
Application Centric Infrastructure (ACI) 243
application development. *See also* application security
application provisioning and deprovisioning 260
automation and scripting 278–279
diversity 278
elasticity 279–280
integrity measurement 261
Open Web Application Security Project (OWASP) 204, 276–277
programming testing methods
compile-time errors 266–267
fuzz testing 269–270
input validation 80, 267–268
penetration testing 266
runtime errors 266–267
static and dynamic code analysis 269
stress testing 80, 266
scalability 279–280
secure coding 261–263
software development environments 257–260
software development lifecycle (SDLC) 78, 261–262, 263–265, 468, 868
vulnerabilities and attacks 74–75
API attacks 55, 85–86, 602
backdoors 149, 271, 275
buffer overflows 75–76, 77, 149, 271–272, 275
code injection 149, 273–274, 276
cross-site request forgery (XSRF) 149, 272, 275
cross-site scripting (XSS) 54, 68–70, 110, 149, 272, 275, 601
directory traversal 75–76, 149, 274–275, 276
DLL injection 74
driver manipulation 89
error handling 79–82
LDAP injection 74
memory/buffer 77–78, 88, 149, 271–272, 275
pass the hash 89–90
pointer dereferencing 75–76
privilege escalation 67–68, 201, 770
race conditions 79
remote code execution (RCE) 78, 146, 149, 275
replay 82–85
request forgeries 85–86
resource exhaustion 87–88

SQL injection (SQLi) 54, 70–74, 273–274
SSL stripping 88–89
summary of 275–276
XML injection 74–75
zero-day attack 149, 275, 276
application logs 792–793
application management, mobile 576–578
application programming interfaces. *See* APIs (application programming interfaces)
application scanners 182
application security 463–464, 475–476, 612. *See also* application development
allow lists 467, 578, 583, 822
application shielding 471
authentication 298
block/deny lists 467–468, 822–823
code signing 466–467
disk encryption 473
dynamic code analysis 470–471
fuzzing 471
hardening 471
hardware root of trust 476–477
Hypertext Transfer Protocol (HTTP) 436–437, 465–466, 577
input validation 464
manual code review 470
mobile devices 581
open ports/services 471–472
operating system 473–474
patch management 474–475
registry 472
sandboxing 452, 478–479
secure coding practices 468
secure cookies 465
self-encrypting drives (SEDs) 475–476
static code analysis 468–469
Trusted Platform Module (TPM) 477–478
whitelisting 578, 583
application service providers (ASPs) 139, 231
application-aware devices 518
application-based segmentation 489–490
application-level gateways (ALGs) 529
approved lists 822
AppScan 204
APs (access points)
rogue 99
security 562–563
APT29 (Cozy Bear) 346
apt-get install snmp snmpwalk command 436
APTs (advanced persistent threats) 120–121, 451, 770
archive.org 147
Arduino 340
ARF (Asset Reporting Format) 885
ARO (annualized rate of occurrence) 922
ARP (Address Resolution Protocol)
poisoning 105, 722
spoofing 513
arp command 721–722
artifacts, forensic 853
artificial intelligence and machine learning (AI/ML) 50–51, 788
As-if Infinitely Ranged (AIR) integer model 77
ASLR (address space layout randomization) 76, 265, 272
ASPs (application service providers) 139, 231
assertion parties (SAML) 659
assertions 623
assessments, security. *See* security assessments
Asset Identification (AI) 885, 941–942
asset management 909–910
Asset Reporting Format (ARF) 885
asset values 921, 922
asymmetric encryption 411–413
ATC (air traffic control) 349–350

ATT&CK framework (MITRE) 18, 128–129, 176, 205, 223, 767–768
Attack Complexity (AC) metric 183
Attack Vector (AV) metric 183
attestation 294, 460–461
attribute-based access control (ABAC) 638–645, 678, 679
audio steganography 415–416
auditing 635, 639
  audit logs 869–870
  audit trails 870
  cloud computing 604, 609
auditors 947
AUPs (acceptable use policies) 898, 900
Australia Privacy Principles (APP) 220
  802.1X standard 510, 553–556, 562, 664–667, 673
  AAA framework 304–306
  attestation 294
  authenticated encryption (AE) 404
  authenticated modes 404
  authentication applications 298
  biometric systems 300, 378, 625–626, 869
    crossover error rate (CER) 304
    efficacy of 302
    errors with 626
    false acceptance rate (FAR) 303, 626
    false rejection rate (FRR) 303, 626
    fingerprints 300–301
    gait analysis 302
    iris recognition 301
    retina scanning 301
    vein authentication 302
    voice/speech recognition 302
  captive portals 559
  Challenge-Handshake Authentication Protocol 673
  challenge-response authentication (CRA) 571–572
  by characteristic attributes 625–626
  CIA (confidentiality, integrity, availability) 289
  cloud versus on-premises requirements 306–307
  context-aware authentication 658
  definition of 289–291, 625
  directory services 291–292
  embedded systems 363
  Extensible Authentication Protocol (EAP) 553–556, 664–667
    EAP-FAST 556, 666
    EAP-MD5 556, 666
    EAP-TLS 556, 666
    EAP-TTLS 556, 666
    LEAP 666
    PEAP 556, 666
  federation 292–293, 556–557, 658
  hardware security modules (HSMs) 656
  HMAC-based one-time password (HOTP) 295–296
  Kerberos 82–83, 89, 292, 553, 668–670, 673
  by knowledge 625, 656–657
  Lightweight Directory Access Protocol (LDAP) 291, 442, 667–670
    injection attacks 74, 144
    Lightweight Directory Access Protocol over SSL (LDAPS) 432
  logs 789–796
  multifactor 304–306, 657
  mutual 668–670
  OAuth 661–662
  OpenID and OpenID Connect 663–664
  by ownership 625
  phone call 299–300
  push notifications 299
  remote
    Challenge-Handshake Authentication Protocol (CHAP) 670–672, 673

RADIUS 556–557, 672–673
Remote Access Service (RAS) 670–672
TACACS+ 672–673
Security Assertion Markup Language (SAML) 659–661
Short Message Service (SMS) 296–297
single sign-on (SSO) 292, 373, 624, 658–659
smart card 299–300, 629
static codes 298
summary of 673
time-based one-time password (TOTP) 295
token key 297
Trusted Platform Module (TPM) 294, 655
two-factor 298
Wi-Fi Protected Setup (WPS) 558–559
authentication attacks 55, 602
Authentication Header (AH) 437, 520
authentication servers 555, 665
authenticators 555, 665
Author Domain Signing Practices (ADSP) 110
authorization 290, 306
authorized hackers 121
Auto (Android) 347
automated indicator sharing (AIS) 125
automation
application development 278–279
auto-updates 474–475
facility 345
autonomous underwater vehicles (AUVs) 353–354
Autopsy 747, 850
AUVs (autonomous underwater vehicles) 353–354
availability 289
resource exhaustion 87–88
restoration order 330–331
site resiliency and 221–222
Availability Impact (I) metric 184
avalanche effect 463
avoidance, risk 918
awareness, risk 921
AWS (Amazon Web Services) 244, 603, 853
Azure 232–233, 603, 853

B

backdoors 42–43, 149, 271, 275
background checks 899
backups 158
cloud 326
comparison of 326–327
copy 326
differential 326, 328
disk 326
full 326, 328–331
image 326
incremental 326, 328
NAS (network-attached storage) 326
offsite 327
online versus offline 326
snapshot 326
tape 326
badges 373, 382
baiting 19
balancers, load 319–320
bandwidth monitors 804
barricades 370–371
base groups 182
baseband radio 359
baselining 213, 539–542
Bash 113
Basic Encoding Rules (BER) 697
basic input/output system (BIOS) 851
BCDR (business continuity and disaster recovery) 139, 232
BCPs (business continuity plans) 773–774, 929
beamforming 560
Bell-LaPadula 677

benchmarks 885–888
BER (Basic Encoding Rules) 697
BGP (Border Gateway Protocol)
  hijacking 535–536
BIA (business impact analysis) 773, 926–927
Biba 677
binaries 278
binary planting 74
biometric systems 300, 378, 625–626, 869
  crossover error rate (CER) 304
  efficacy of 302
  errors with 626
  false acceptance rate (FAR) 303, 626
  false rejection rate (FRR) 303, 626
  fingerprints 300–301
  gait analysis 302
  iris recognition 301
  retina scanning 301
  vein authentication 302
  voice/speech recognition 302
BIOS (basic input/output system) 851
birthday attacks 56
BiSL (Business Information Services Library) 882
Bitcoin-related SMS scams 12
BitTorrent 529
black hat hackers 121
black-box testing 80
blackhole DNS servers 223
Blackhole exploit kit 44, 111–112
blacklisting 578, 583
blanket purchase agreements (BPAs) 903
blind hijacking 84
blind SQL injection 73
block all. *See* implicit deny
block ciphers 411
blockchain 409–410
block/deny lists 467–468, 578, 583, 822–823
blocking 417
Blowfish 412
blue teams 205, 902
Bluetooth 570–571
  bluejacking 100, 570–571
  bluesnarfing 99–100, 570–571
bollards 370–371
Boolean technique 74
boot integrity
  boot attestation 460–461
  definition of 458–459
  measured boot 459–460
  Unified Extensible Firmware Interface (UEFI) 459
Border Gateway Protocol (BGP)
  hijacking 535–536
bots and botnets 37–38, 111–112, 580
BPAs. *See* blanket purchase agreements (BPAs); business partnership agreements (BPAs)
BPDU (Bridge Protocol Data Unit)
  guard 512
bring-your-own-device (BYOD) 215, 572, 574–576, 581, 588–590, 826, 898
broadcast storm prevention 512
  BPDU guard 512
  DHCP snooping 512–513
  loop protection 512
  MAC filtering 513
brute-force attacks 45, 749
buckets 605
buffer overflows 75–76, 77, 149, 271–272, 275, 522
bug bounties 202–203
BugCrowd 203
building loss 925
burning 386
Burp Suite Professional 204
buses, controller area network (CAN) 347–348
business continuity and disaster recovery (BCDR) 139, 232
business continuity plans (BCPs) 773–774, 929

business impact analysis (BIA) 773, 926–927
business partnership agreements (BPAs) 903
BYOD. *See* bring-your-own-device (BYOD)

C

cables
  locks 379
  malicious USB 48
CAC (Common Access Card) 629
cache
  ARP cache poisoning 105
  caching proxy 514
  DNS cache poisoning 108–110
  forensic acquisition 852
Cain and Abel 44
California Consumer Privacy Act (CCPA) 214, 220, 880
call management systems (CMSs) 351
Call Manager log files 799–800
CAM (content addressable memory) 106
cameras
  centralized versus decentralized 375
  closed-circuit television (CCTV) 376–377, 870
  motion recognition 376
  object detection 376
camouflage 265, 377
CAN (controller area network) bus 347–348
Canada, Personal Information Protection and Electronic Data Act (PIPEDA) 220
Canonical Encoding Rules (CER) 697
capital expenditure (CapEx) 598
captive portals 559
capture, packet. *See* packet capture and replay
capture the flag 902
card cloning attacks 48–49
CarPlay, Apple 347
carrier unlocking 584
CAs (certificate authorities) 466, 556, 689–691, 829
CASBs (cloud access security brokers) 142–143, 611–612, 614
cat command 734–735
CBC (Cipher Block Chaining) mode 405
CBT (computer-based training) 901
CBWFQ (class-based weighted fair queuing) 536
CCE (Common Configuration Enumeration) 886
CCleaner 51
CCPA (California Consumer Privacy Act) 214, 220, 880
CCSS (Common Configuration Scoring System) 886
CCTV (closed-circuit television) 376–377, 870
CD (continuous delivery) 279
CDP (clean desk policy) 23, 899, 900
Cellebrite 850–851
cellular connection methods and receivers 572–573
Center for Internet Security (CIS) 164, 881, 883
centralized access control 640, 679
centralized cameras 375
centralized controllers 242
CER (Canonical Encoding Rules) 697
CER (crossover error rate) 304, 626
.cer file extension 697
CERT (Community Emergency Response Team) 77
certificate authorities (CAs) 466, 556, 689–691, 829
certificate revocation lists (CRLs) 533, 689–690, 691, 829
certificate signing requests (CSRs) 689

certificates 625, 626–627
  attributes 691–692
  chaining 696
  expiration 693
  formats 697
  pinning 698
  Subject Alternative Name 693
  types of 694–696
  updating/revoking 829–830
CFB (Cipher Feedback) mode 406
chain of custody 789, 844
chain of trust 699
Challenge-Handshake Authentication Protocol (CHAP) 673
challenge-response authentication (CRA) 49–50, 102, 571–572
change management 909
CHAP (Challenge-Handshake Authentication Protocol) 82–83, 670–672, 673
characteristic attributes, authentication by 625–626
Check Point 518
checksums 857, 870
chief information officers (CIOs) 903
chief security officers (CSOs) 930
chkdsk command 157
chmod command 644–645, 736–737
choose-your-own-device (CYOD) 588–590
CI (continuous integration) 279
CIA (confidentiality, integrity, availability) 221, 263, 289
CIDR (classless interdomain routing) netblock 203–204
CIOs (chief information officers) 903
Cipher Block Chaining (CBC) mode 405
Cipher Feedback (CFB) mode 406
cipher suites 409–411
CIRT. *See* incident response (IR) teams
CIS (Center for Internet Security) 164, 881, 883
CISA (Cybersecurity and Infrastructure Security Agency) 353–354
Cisco
  Application Centric Infrastructure (ACI) 243
  Application Policy Infrastructure Controller (APIC) 243
  Cisco Discovery Protocol (CDP) 107
  Email Security Appliance (ESA) 111
  Identity Services Engine (ISE) 590
  Mutiny Fuzzing Framework 269
  NetFlow 187, 525, 809–810
  OpenDNS 509–510
  security advisories and bulletins 179
  Talos 347
  Umbrella 509
Clark-Wilson 677
class-based weighted fair queuing (CBWFQ) 536
classification
  asset 941–942
  data 904–905
classless interdomain routing (CIDR) netblock 203–204
clean desk policy (CDP) 23, 899, 900
clean pipe 112
clickjacking 84
client-based VPNs (virtual private networks) 497
clientless VPNs (virtual private networks) 497, 507–508
clientless web access 507
clients, thin 235–236, 508
client-side execution 267
client-side validation 268
clock, secure 477
cloning
  MAC (media access control) 106
  SIM (subscriber identity module) cards 580, 584
closed-circuit television (CCTV) 376–377, 870

cloud access security brokers (CASBs) 142–143, 611–612, 614
cloud computing
advantages of 138
attacks and vulnerabilities 52–55, 123, 137–143, 601–603
authentication 306–307
backups 326
cloud access security brokers (CASBs) 142–143, 611–612, 614
cloud service providers (CSPs) 139, 233, 598, 853–854
community cloud 140, 233
definition of 138
fog and edge computing 234–235
forensic acquisition 853–854
hybrid cloud 140, 233
managed detection and response (MDR) 234
managed service providers (MSPs) 233–234
models 231–232
off-premises versus on-premises services 234
private cloud 140, 232–233
public cloud 140, 232
resilience 325
security assessments 598
attacks 601–603
threats 598–600
security controls 595, 598
API inspection and integration 607, 610
compute 607, 611
container security 608–609
dynamic resource allocation 607–608, 611
high availability across zones 603, 609
instance awareness 608, 611
integration and auditing 604, 609
native versus third-party 615
network 606–607, 610
resource policies 603, 609
secrets management 604, 609
security groups 607, 611
storage 605, 610
summary of 608–609
virtual private cloud endpoint 608, 611
security solutions
application security 612
cloud access security brokers (CASBs) 611–612, 614
firewalls 613–614, 615
Secure Web Gateway (SWG) 613, 614
summary of 614–615
storage
encryption 605
high availability 606
permissions 605
replication 605
thin clients 235–236
VPCs (virtual private clouds) 607, 608, 611
Cloud Controls Matrix 884
Cloud Security Alliance (CSA) 139, 603, 884
Cloud Service (Google) 603
cloud service providers (CSPs) 139, 233, 598, 853–854
Cloudflare 440
cloudlets 235
Cluster Server 488
CMSs (call management systems) 351
CMSS (Common Misuse Scoring System) 887
COBIT framework 882
code, infrastructure as 241–243
code security 261–263
code camouflage 265
code checking 79, 265
code injection 149, 273–274, 276

code reuse 179, 270
code signing 466–467, 695, 696
dynamic code analysis 470–471
manual code review 470
static code analysis 468–469
cold aisles 386
cold sites 222
collection, log 186
collisions 55–56, 463
command-and-control (C2) servers 37–38, 107
commands. *See individual commands*
comment delimiters 73
Common Access Card (CAC) 629
Common Configuration Enumeration (CCE) 886
Common Configuration Scoring System (CCSS) 886
Common Misuse Scoring System (CMSS) 887
common names (CNs) 692
Common Object Request Broker Architecture (CORBA) 86
Common Platform Enumeration (CPE) 886
Common Remediation Enumeration (CRE) 886
Common Security Advisory Framework (CSAF) 164
Common Vulnerabilities and Exposures (CVEs), Wi-Fi 78, 125, 146, 177, 571, 886
Common Vulnerability Reporting Framework (CVRF) 164
Common Vulnerability Scoring System (CVSS) 182–186, 886
Common Weakness Enumeration (CWE) 75, 886
Common Weakness Scoring System (CWSS) 887
communications
communication plans 771–772
embedded systems
5G 357–358
baseband radio 359
NarrowBand 358
subscriber identity module (SIM) cards 360
Zigbee 360–361
community cloud 140, 233
Community Emergency Response Team (CERT) 77
community ports 491
company policies 878–879
compensating controls 871, 872
compilers 278
compile-time errors 81–82, 266–267
compliance, software 918
computer certificates 696
computer incident response teams. *See* incident response (IR) teams
computer-based training (CBT) 901
Concealment 415
concentrators, VPN 495
conditional access 678, 679
confidence tricks 19
Confidential information 905, 941–942
Confidentiality Impact (C) metric 184
configuration management 164, 213
configuration reviews 182
mitigation techniques 824
certificates, updating/revoking 829–830
content filter/URL filter 828–829
data loss prevention (DLP) 825–826
firewall rules 825
mobile device management (MDM) 825–826
secure configuration guides 885–888
weak configurations 150–155

connection methods and receivers
Bluetooth 570–571
cellular 572–573
Global Positioning System (GPS) 572, 584
near-field communication (NFC) 570–571
Radio frequency identification (RFID) 571–572
satellite communications (SATCOM) 573
secure implementation best practices 573–574
containers 236–240, 608–609
containment, incident response (IR) 763–764, 830–831
content addressable memory (CAM) 106
content filters 533, 828–829
content management 576–578
context-aware authentication 658
continuity of operations plans (COOPs) 774–775, 929
continuous delivery (CD) 279
continuous deployment 279
continuous integration (CI) 279
continuous monitoring 139, 278
continuous validation 278
Control Objectives for Information and Related Technology (COBIT) 882
control systems, diversity in 332
controller area network (CAN) bus 347–348
controller-pilot data link communications (CPDLC) 349–350
controllers 562–563, 946
controls. *See also* physical security
compensating 871, 872
corrective 870, 872
detective 869–870, 872
deterrent 870–871, 872
managerial 868
operational 868, 869
physical 871–872
preventative 869, 872
technical 868, 869
convert command 156
cookie hijacking 465
cookies 465
cookies, secure 465
COOPs (continuity of operations plans) 774–775, 929
Coordinated Universal Time (UTC) 440, 845
COPE (corporate-owned, personally enabled) environments 572, 588
copy backups 326
CORBA (Common Object Request Broker Architecture) 86
corporate incidents 775
corporate-owned, personally enabled (COPE) environments 572, 588–590
corrective controls 870, 872
correlation, log 186
correlation, Security Information and Event Management (SIEM) 788–789
Counter (CTR) mode 404, 408–409
counterintelligence 860
Counter-mode/CBC-MAC protocol (CCMP) 552
counters, secure 477
county names, certificate 692
cover-files 416
Cozy Bear 346
CPE (Common Platform Enumeration) 886
CRA (challenge-response authentication) 49–50, 102, 571–572
cracking passwords 46

CRE (Common Remediation Enumeration) 886
CREATE DATABASE statement 70
CREATE INDEX statement 71
CREATE TABLE statement 71
credentials
  credentialed vulnerability scans 182, 349–350
  harvesting 18
  policies 906–908
crimeware 44
criminal syndicates 120
Critical information 942
critical systems, identification of 929
CRLs (certificate revocation lists) 533, 689–690, 691, 829
crossover error rate (CER) 304, 626
cross-site request forgery (XSRF) 85–86, 149, 272, 275, 602
cross-site scripting (XSS) 54, 68–70, 110, 149, 272, 275, 464, 601
.crt file extension 697
cryptography 396. *See also* encryption; hashing; secure protocols
  algorithms 498
  blockchain 409–410
  cipher suites 409–411
  common use cases 417–418
  cryptographic attacks
    birthday 56
    collision 55–56
  cryptographic protocols 551
    Advanced Encryption Standard (AES) 552
    Counter-mode/CBC-MAC protocol (CCMP) 552
    Simultaneous Authentication of Equals (SAE) 551, 552
    summary of 552
    Wi-Fi Protected Access 2 (WPA2) 551
    Wi-Fi Protected Access 3 (WPA3) 551–552
  definition of 391
  digital signatures 395–396, 520
  diversity in 331
  elliptic-curve cryptography (ECC) 399–400
  encryption 159, 362
    cloud computing 605, 610
    data at rest 218
    data in transit/motion 218
    data in use/processing 218
    disk 473
    entropy 419
    homomorphic 417
    international mobile subscriber identity (IMSI) 49, 358, 584
    mobile device management (MDM) 578–580
    symmetric/asymmetric 411–413
    vulnerabilities 150–151
  entropy 419
  keys
    ephemeral 403
    key exchanges 399
    key signing keys (KSKs) 427
    length of 396
    password 655
    personal unblocking keys (PUKs) 360
    public/private 436–437
    Secure Shell (SSH) 625, 628
    stretching 397
    zone signing keys (ZSKs) 427
  lightweight 414–415
  limitations of 418–420
  modes of operation 403–409
    authenticated 404
    Cipher Block Chaining (CBC) 405
    Cipher Feedback (CFB) 406
    counter 404
    Counter (CTR) 408–409
    Electronic Code Book (ECB) 404
    Output Feedback (OFB) 407
    unauthenticated 404

perfect forward secrecy 400–401
post-quantum 402
Public Key Cryptography Standards (PKCS) 412
quantum 401–402
communications 401–402
computing 402
definition of 401
salting 397–398, 462–463
steganography 415
audio 415–416
homomorphic 417
image 416–417
video 416
cryptomalware 33–34
CSA (Cloud Security Alliance) 139, 603, 884
CSAF (Common Security Advisory Framework) 164
CSF (Cybersecurity Framework) 882, 884
CSIRT. *See* incident response (IR) teams
CSOs (chief security officers) 930
CSPs (cloud service providers) 139, 233, 598, 853–854
CSRF (cross-site request forgery) 602
CSRs (certificate signing requests) 689
CTR (Counter) mode 404, 408–409
Cuckoo 731–732
curl command 724–725
custodians, data 946
custody, chain of 789, 844
CVE (Common Vulnerability and Exposure) 78, 125, 146, 177, 886
CVE Numbering Authorities (CNAs) 179
CVRF (Common Vulnerability Reporting Framework) 164
CVSS (Common Vulnerability Scoring System) 182–186, 886
CWE (Common Weakness Enumeration) 75, 886
CWSS (Common Weakness Scoring System) 887
cyber kill chain 770–771
Cybersecurity and Infrastructure Security Agency (CISA) 353–354
Cybersecurity Framework (CSF) 882, 884
cybersecurity insurance 918
cybersecurity resilience. *See* resilience
CYOD (choose-your-own-device) 588

D

DAC (discretionary access control) 674–676, 679
DAEAD (deterministic authenticated encryption with associated data) 404
DAI (Dynamic ARP Inspection) 105
dark web 124–125, 143
Darkleech 146–147
dashboards, SIEM 786–789
DAST (dynamic application security testing) 470–471
data at rest 156, 218
data blockers, USB 379–380
data breaches
data types and asset classification 941–942
fines 940
identity theft 940
impact assessment 948
information lifecycle 947–948
intellectual property theft 940
notifications of 855–856, 941
personally identifiable information (PII) 943
privacy enhancing technologies 944–945
privacy notices 949
protected health information (PHI) 944
reputation damage from 940
response and recovery controls 220–221

security roles and responsibilities 945–947
terms of agreement 948
data classification 904–905
data controllers 946
data custodians/stewards 946
data destruction, secure 386–387
Data Encryption Standard (DES) 412
data exfiltration 907–908
data exposure 267
data governance 904–905
data in transit/motion 156, 218
data in use/processing 156, 218
data input 186
data labeling 676
data loss prevention (DLP) 139, 214–215, 453, 582, 586, 699, 825–826, 871
data masking 216–218, 945
data minimization 944–945
data owners 946
data privacy. *See* privacy breaches
data privacy officers (DPOs) 905
data processors 946
data protection 214–215
data protection officers (DPOs) 947
data recovery 859
data retention policies 775–776, 906
data sanitization 748–749
data sources
bandwidth monitors 804
Internet Protocol Flow Information Export (IPFIX) 811–813
log files 789
application 792–793
authentication 789–796
Call Manager 799–800
Domain Name System (DNS) 795–796
dump files 797
journalctl 802
network 790
NXLog 803–804
security 793
Session Initiation Protocol (SIP) 800
syslog/rsyslog/syslog-ng 800–801
system 791–792
Voice over Internet Protocol (VoIP) 799–800
web server 794
metadata 805–806
in email 808
in files 809
on mobile devices 808
on web pages 808–809
NetFlow 809–810
protocol analyzers 813
Security Information and Event Management (SIEM)
alerts 788
correlation 788–789
dashboards 786–789
sensitivity 788
sensors 787
trends 788
sFlow 810–811
vulnerability scan output 785–786
data sovereignty 214–215
data types 941–942
databases 461–462
DC (direct current) 380
DCOM (Distributed Component Object Model) 86
DCS (distributed control systems) 343
DCT (Discrete Cosine Transforms) 417
dd utility 744–745
DDoS (distributed denial-of-service) attacks 37–38, 54, 111–113, 601
dead box forensic collection 858
dead code 270
Dead Peer Detection (DPD) 501
deauthentication attacks 101
decentralized access control 640, 679
decentralized cameras 375

decentralized trust models 698
deception and disruption techniques
  fake telemetry 223
  honeyfiles 223
  honeypots 221–223
DeepSound 415
defense in depth 264
defrag command 158
defragmentation 158
degaussing 387
delegation of access 662
DELETE statement 70
delivery
  continuous 279
  malware 43–45
demilitarized zones (DMZs) 384, 491
denial-of-service (DoS) attacks 88, 122, 267, 601, 770
deny lists 467–468, 578, 583, 822–823
Department of Defense (DoD) security standards 674
deployment, continuous 279
deprovisioning, application 260
DER (Distinguished Encoding Rules) 697
dereferencing, pointer 75–76
DES (Data Encryption Standard) 412
design constraints, embedded systems 361
  authentication 363
  compute 361–362
  cost 363
  crypto 362
  implied trust 363
  inability to patch 362
  network 362
  power 361
  range 363
destruction and disposal services 387
detective controls 869–870, 872
deterministic authenticated encryption with associated data (DAEAD) 404
deterrent controls 869, 870–871, 872
development environments 257–260
development lifecycle. *See* software development lifecycle (SDLC)
devices, forensic acquisition 850–851
devices, mobile. *See* mobile solutions
DevOps 259, 263–265, 278–279
DevSecOps 259, 278–279
DFIR (Digital Forensics and Incident Response) 744
DHCP (Dynamic Host Configuration Protocol) 443
  snooping 512–513
  starvation attack 513
diagrams, configuration 213
Diamond Model of Intrusion Analysis 768–770
dictionary attacks 45, 749
differential backups 326, 328
Diffie-Hellman 500–501
dig command 709–710
DigiCert 691
digital forensics. *See* forensics, digital
Digital Millennium Copyright Act 220
digital rights management (DRM) 67, 219–220
digital signal processors (DSPs) 359
Digital Signature Algorithm (DSA) 396, 412
digital signatures 395–396, 520
digital video recorders (DVRs) 376–377
direct current (DC) 380
directory services 291–292, 442
directory traversal 75–76, 149, 274–275, 276
disablement 635, 639
disassociation attacks 101
disaster analysis 924–925
disaster recovery plans (DRPs) 330–331, 772–773, 926, 928–930
disclosures, public 940
discovery of identity 623–624

discovery tools
  definition of 707
  dig 709–710
  hping 717
  ifconfig 710–711
  ipconfig 710
  netcat 720–721
  netstat 718–720
  nmap 711–714
  nslookup 709–710
  pathping 716–717
  ping 714–716
  ping6 716
  tracert/traceroute 707–709
Discrete Cosine Transforms (DCT) 417
discretionary access control (DAC) 674–676, 679
Disk Cleanup 157
Disk Defragmenter 158
disks
  backups 326
  encryption 473
  forensic acquisition of 848
  hardening 157–159
  redundancy
    definition of 315–316
    multipath 319
    Redundant Array of Inexpensive Disks (RAID) 315–316
  Redundant Array of Inexpensive Disks (RAID) 869
  self-encrypting 475–476
Distinguished Encoding Rules (DER) 697
Distributed Component Object Model (DCOM) 86
distributed control systems (DCS) 343
distributed denial-of-service (DDoS) attacks 37–38, 54, 111–113, 601
Distributed Ledger Technology (DLT) 409
diversity 278, 331–332
DKIM (Domain Keys Identified Mail) 110, 426
DLL (dynamic link library) injection 74, 274
DLP (data loss prevention) 139, 214–215, 453, 582, 586, 699, 825–826, 871
DLT (Distributed Ledger Technology) 409
DMARC (Domain-based Message Authentication, Reporting & Conformance) 111
DMSSEC (Domain Name System Security Extensions) 796
DMZs (demilitarized zones) 384, 491
DNS (Domain Name System) 442–443
  attacks 54
    cloud-based 601
    DDoS (distributed denial-of-service) 37–38, 54, 111–113, 601
    DNS amplification attack 112
    DNS poisoning 108–110, 223
    domain hijacking 108
    domain name kiting 109–110
    domain reputation 110–111
    prevalence of 107
    URL redirection attacks 110
  DNS Security Extensions (DNSSEC) 108, 426–427
  DNS sinkholes 223
  logs 795–796
  OpenDNS 509–510
dnsenum 728–729, 796
DNSSEC (Domain Name System Security Extensions) 108, 426–427, 442–443
Docker 237–240
docker images command 237
docker ps command 238
docker search command 239
Document Object Model (DOM) 68–69

documentation, forensic
admissibility of 843
chain of custody 844
event logs 845–846
interviews 846–847
legal hold 842
reports 846
tagging of 845–846
timelines and sequence of events 844–845
time offset 844
timestamps 844
video 842–843
DOM (Document Object Model) 68–69
Domain Keys Identified Mail (DKIM) 110, 426
domain name kiting 109–110
domain name resolution 442–443
Domain Name System. *See* DNS (Domain Name System)
domain reputation 110–111
domain validation (DV) certificates 694
Domain-based Message Authentication, Reporting & Conformance (DMARC) 111
DoS (denial-of-service) attacks 88, 601, 770
DPD (Dead Peer Detection) 501
DPOs (data privacy officers) 905
Dragonfly 101, 552
driver manipulation 89
drives. *See* disks
DRM (digital rights management) 67, 219–220
drones 205, 353–354, 382–383
DROP INDEX statement 71
DROP TABLE statement 71
DRPs (disaster recovery plans) 772–773, 926, 928–930
DSA (Digital Signature Algorithm) 396, 412
DSPs (digital signal processors) 359
DTP (Dynamic Trunking Protocol) 106
dual parity, striping with (RAID) 316, 318
dual power supplies 321
dual supply power 321–322
due care 900
due diligence 900
due process 900
dump files 797
dumpster diving 13
duties, separation of 898, 900
DV (domain validation) certificates 694
DVRs (digital video recorders) 376–377
dynamic application security testing (DAST) 470–471
Dynamic ARP Inspection (DAI) 105
dynamic code analysis 269, 470–471
Dynamic Host Configuration Protocol. *See* DHCP (Dynamic Host Configuration Protocol)
dynamic link library (DLL) injection 74, 274
dynamic resource allocation 607–608, 611
Dynamic Trunking Protocol (DTP) 106

E

EAP (Extensible Authentication Protocol) 553–556, 664–667
EAP-FAST 556, 666
EAP-MD5 556, 666
EAP-TLS 556, 666
EAP-TTLS 556, 666
LEAP 666
PEAP 556, 666
Easter eggs 39–40
east-west traffic 492
ECB (Electronic Code Book) 404
ECC (elliptic-curve cryptography) 399–400
ECDSA (Elliptic Curve Digital Signature Algorithm) 551–552
edge computing 234–235
e-discovery 858–859

EDR (endpoint detection and response) 452–453
education, user 22–24, 899, 901–902
EEA (European Economic Area) 214, 220
EER (equal error rate). *See* crossover error rate (CER)
eEye Digital Security, Retina Web Security Scanner 204
EFS (Encrypting File System) 694
EIGamal 412
elasticity 279–280
electrical metallic tubing (EMT) 385
electromagnetic (EM) frequency band 102
Electronic Code Book (ECB) 404
electronic locks 379
electronic serial numbers (ESNs) 49, 584
eliciting information 15–16
Elliptic Curve Digital Signature Algorithm (ECDSA) 551–552
elliptic-curve cryptography (ECC) 399–400
elliptic-curve techniques 412
EM (electromagnetic) frequency band 102
email
  attack vectors 122
  certificates 696
  email protocol port numbers 441
  email servers 145
  metadata in 808
  Spam 13
  SPIM (Spam over Internet Messaging) 13
  synchronization 440
Email Security Appliance (ESA) 111
embedded systems
  aircraft 348–350
  Arduino 340
  communication considerations
    5G 357–358
    baseband radio 359
    NarrowBand 358
    subscriber identity module (SIM) cards 360
    Zigbee 360–361
  constraints 361
    authentication 363
    compute 361–362
    cost 363
    crypto 362
    implied trust 363
    inability to patch 362
    network 362
    power 361
    range 363
  definition of 339
  drones 353–354
  Field-Programmable Gate Array (FPGA) 340
  heating, ventilation, and air conditioning (HVAC) 352–353
  industrial control systems (ICSs) 341–343
  Internet of Things (IoT) 38, 98, 113, 344–346, 358, 414
  medical systems 347
  multifunction printers (MFPs) 354
  Raspberry Pi 339
  real-time operating systems (RTOSs) 355
  smart meters 350
  supervisory control and data acquisition (SCADA) 341–343
  surveillance systems 355–356
  system on a chip (SoC) 356–357
  vehicles 347–348
  Voice over Internet Protocol (VoIP) 350, 799–800
emergency preparedness logs 383
EMT (electrical metallic tubing) 385
Encapsulating Security Payload (ESP) 437, 503, 520
EnCase 850–851

Encrypting File System (EFS) 694
encryption 159, 362
  cloud computing 605, 610
  data at rest 218
  data in transit/motion 218
  data in use/processing 218
  disk 473
  entropy 419
  homomorphic 417
  international mobile subscriber identity (IMSI) 49, 358, 584
  mobile device management (MDM) 578–580
  modes of operation 403–409
    authenticated 404
    Cipher Block Chaining (CBC) 405
    Cipher Feedback (CFB) 406
    Counter (CTR) 404, 408–409
    Electronic Code Book (ECB) 404
    Output Feedback (OFB) 407
    unauthenticated 404
  symmetric/asymmetric 411–413
  vulnerabilities 150–151
end of life (EOL) 904
end of service life (EOSL) 904
end users 947
endpoint detection and response (EDR) 452–453
endpoint DLP systems 214
endpoint protection 451
endpoint security solutions 822
  approved lists 822
  block/deny lists 467–468, 578, 583, 822–823
  quarantine 823–824
end-to-end headers (HTTP) 466
energy management, SCADA control systems 342–343
engagement, rules of 200
enterprise environments
  API considerations 216
  configuration management 213, 215–216
  data masking 216–218
  data protection 214–215
  data sovereignty 214–215
  deception and disruption techniques
    fake telemetry 223
    honeyfiles 223
    honeypots 221–223
  digital rights management (DRM) 219–220
  DNS sinkholes 223
  encryption 218
  hashing 218–219
  response and recovery controls 220–221
  site resiliency 221–222
enterprise resource planning (ERP) 883
entropy 419
enumerations 886
env command 739
environmental disaster 924
environmental groups 182
environmental variables 740
environments, software development 257–260
  known 198
  partially known 199
  unknown 198–199
EOL (end of life) 904
EOSL (end of service life) 904
ephemeral keys 403
equal error rate. *See* crossover error rate (CER)
eradication phase, incident response (IR) 764
ERP (enterprise resource planning) 883
error handling 79–82
  compile-time errors 81–82, 266–267
  error-based technique 74
  input handling 80
  runtime errors 81–82, 266–267

escalation, privilege 67–68, 201, 941
escape attacks, VM (virtual machine) 248–249
escrow, key 699
ESNs (electronic serial numbers) 49, 584
ESP (Encapsulating Security Payload) 437, 503, 520
ethical hacking. *See* penetration testing
ETSI (European Telecommunications Standards Institute) 235
EU (European Union)
  European Economic Area (EEA) 214, 220
  European Telecommunications Standards Institute (ETSI) 235
  General Data Protection Regulation (GDPR) 42, 214, 220, 356, 434, 453, 760, 855, 878–879, 947
  Information Society Directive 220
EV (extended validation) certificates 694
event logs 845–846
Event Viewer 791–792
events, sequence of 844–845
  time offset 844
  timestamps 844
evidence, forensic
  acquisition 847–854
    artifacts 853
    cache 852
    checksums 857
    data breach notification laws 855–856
    definition of 847
    device 850–851
    disk 848
    firmware 851
    hashing 856–857
    integrity 856
    network 852–853
    operating system 850
    order of volatility 848
    on-premises versus cloud 853–854
    random-access memory (RAM) 848–849
    regulatory and jurisdictional 855
    right-to-audit clauses 854
    snapshot 851–852
    swap/pagefile 849–850
  admissibility of 843
  chain of custody 844
  e-discovery 858–859
  event logs 845–846
  interviews 846–847
  legal hold 842
  preservation 858
  provenance 857–858
  reports 846
  tagging of 845–846
  timelines and sequence of events 844–845
    time offset 844
    timestamps 844
  video 842–843
evil twin attacks 98–99
exam preparation
  final review and study 953–954
  hands-on activities 953
  Pearson Test Prep practice test 954
  test lab, building 953
exam updates 02.0004–02.0026
exchanges, key 399
executives, security roles and responsibilities 945–947
exercises
  simulations 766–767
  tabletop 765–766
  walkthrough 766
exFAT 850
exfiltration 770, 907–908
expiration, certificates 693
explicit allow/deny 528
Exploit code maturity (E) metric 184
exploit kits 44
Exploitability metrics 183–184

exploitation frameworks 747–748, 770
Extended Detection and Response (XDR) 189
extended validation (EV) certificates 694
Extensible Authentication Protocol. *See* EAP (Extensible Authentication Protocol)
Extensible Configuration Checklist Description Format (XCCDF) 885
Extensible Markup Language (XML) injection 74–75, 273–274
external actors 122
external risk 917. *See also* risk management
extinguishers, fire 381
extranets 492–493, 899

F

f8-mode (SRTP) 430
FAA (Federal Aviation Administration) 348–349, 353, 382–383
facility automation 345
facility codes 373
fail-closed 927
fail-open 927
failure, single point of 156, 926
failure in time (FIT) 926
fake telemetry 223
false acceptance rate (FAR) 303, 626
false negatives 181, 519, 520
false positives 181, 518, 520
false rejection rate (FRR) 303, 626
Faraday cages 383, 562–563
FAST (Flexible Authentication via Secure Tunneling) 556
Fast Identity Online (FIDO) 297
FAT 850
FDE (full-disk encryption) 473, 475–476
fdisk -l command 157
FEAT command 433
Federal Aviation Administration (FAA) 348–349, 353, 382–383
Federal Information Security Management Act (FISMA) 776
Federal Risk and Authorization Management Program (FedRAMP) 599
Federal Trade Commission (FTC) 17, 221
federated identity management (FIM) 658
federation 292–293, 623–624, 658, 672
FedRAMP (Federal Risk and Authorization Management Program) 599
fencing 380–381
FFmpeg 416
FIDO (Fast Identity Online) 297
Field-Programmable Gate Array (FPGA) 340
file and code repositories 127
file integrity monitors 542
file manipulation 732–733
  cat command 734–735
  chmod command 736–737
  grep command 735–736
  head command 733
  logger command 737–738
  tail command 734
file servers 144
file transfer 440
File Transfer Protocol. *See* FTP (File Transfer Protocol)
fileless viruses 37
files
  log 789
    application 792–793
    authentication 789–796
    Call Manager 799–800
    Domain Name System (DNS) 795–796
    dump files 797

journalctl 802
network 790
NXLog 803–804
security 793
Session Initiation Protocol (SIP) 800
syslog/rsyslog/syslog-ng 800–801
system 791–792
Voice over Internet Protocol (VoIP) 799–800
web server 794
metadata in 809
filtering
content/URL 533, 828–829
MAC (media access control) 513
packet 528
financial information. *See* personally identifiable information (PII)
Financial Services Information Sharing and Analysis Center (FS-ISAC) 124
fines 940
fingerprint authentication 300–301
fire
disaster analysis for 924–925
suppression 381
firewalls 146, 198
appliance 534
application-level gateway (ALG) 529
in cloud 613–614, 615
configuration 529–533
content URL/filtering 533
hardware versus software 534
host-based 457–458, 534
multihomed connections 532
NAT gateway 529
network-based application layer 530
next-generation firewall (NGFW) 453–454, 524
packet filtering 528
personal 534
purpose of 526–528
rules 528, 825
unified threat management (UTM) 524
virtual 534–535
web application 531
wireless security 562
firmware
firmware over-the-air (OTA) updates 583
forensic acquisition of 851
FIRST (Forum of Incident Response and Security Teams) 180
FISMA (Federal Information Security Management Act) 776
FIT (failure in time) 926
flash drives, malicious 47–48
Flexible Authentication via Secure Tunneling (FAST) 556
flood, disaster analysis for 925
flooding, MAC (media access control) 106
FM200 381
fog computing 234–235
footprinting 205
Forcepoint 533
Forefront Identity Manager 658
Foremost 415
Forensic Toolkit (FTK) 747, 850–851
forensics, digital
acquisition
artifacts 853
cache 852
checksums 857
data breach notification laws 855–856
definition of 847
device 850–851
disk 848
firmware 851
hashing 856–857
integrity 856
network 852–853
operating system 850
order of volatility 848

on-premises versus cloud 853–854
random-access memory (RAM) 848–849
regulatory and jurisdictional 855
right-to-audit clauses 854
snapshot 851–852
swap/pagefile 849–850
data recovery 859
definition of 744, 837
Digital Forensics and Incident Response (DFIR) 744
documentation/evidence
admissibility of 843
chain of custody 844
event logs 846
interviews 846–847
legal hold 842
reports 846
tagging of 845–846
timelines and sequence of events 844–845
video 842–843
e-discovery 858–859
nonrepudiation 859–860
preservation 858
provenance 857–858
strategic intelligence/counterintelligence 860
tools
Autopsy 747
dd 744–745
FTK Imager 747
memdump 745
WinHex 746
forgeries, request 85–86
formats, certificate 697
Forum of Incident Response and Security Teams (FIRST) 180
forward proxy 516
forward secrecy 400–401
FPGA (Field-Programmable Gate Array) 340
frameworks
exploitation 747–748
IT security 881–884
FreeBSD 676
frequency distributions 159
FRR (false rejection rate) 303, 626
fsck command 158
FTC (Federal Trade Commission) 17, 221
FTK (Forensic Toolkit) 747, 850–851
FTP (File Transfer Protocol)
FTP servers 147–148
FTPS (File Transfer Protocol, Secure) 432–433
SFTP (Secure File Transfer Protocol) 434
full backups 326, 328–331
full tunnel mode, SSL/TLS VPN 508
full-disk encryption (FDE) 473, 475–476
functions, hash 218–219
fuzz testing 80, 269–270, 471
fuzzers 269–270

G

gait analysis 302
Galois Message Authentication Code (GMAC), AES in 498
Galois/Counter Mode (GCM) 498, 551–552
gamification 902
gapping 384
gateways
application-level 529
NAT 529
transit 246–247
GCM (Galois/Counter Mode) 498, 551–552
GDPR (General Data Protection Regulation) 42, 214, 220, 356, 434, 453, 760, 855, 878–879, 947
general-purpose I/O GPIO framework extension (GpioClx) 477

generators 321
generic accounts 629
Generic Routing Encapsulation (GRE) 520
geofencing 572–573, 578–580
geographic dispersal 315
geolocation 578–580, 639
geotagging 572–573, 584, 586, 639
GitHub repositories 8, 18, 127, 203, 258
GitLab 127, 258
Global Positioning System (GPS) 572, 584
Global Regular Expression Print (grep) 735–736
GMAC (Galois Message Authentication Code), AES in 498
Gnutella 530
Golden SAML attacks 293
Google
  Cloud 233, 603, 853
  Google Pay 584
  Kubernetes 239–240
  OAuth 2.0 292
  Secret Manager 604
governance, risk, and compliance (GRC) 880, 904–905
GpioGlx (general-purpose I/O GPIO framework extension) 477
GPOs (group policy objects) 474
GPS (Global Positioning System) 572, 584
Gramm-Leach-Bliley (GLB) Act 880
GraphQL 86
gray hat hackers 121
gray-box testing 80
GRC (governance, risk, and compliance) 880, 904–905
GRE (Generic Routing Encapsulation) 520
grep command 735–736
group policy objects (GPOs) 474
groups
  base 182
  environmental 182
  security 607, 611
  temporal 182
Grover's algorithm 402
guards 377
guest accounts 629
Guidelines for Evidence Collection and Archiving 848

H

HA (high availability) 329–330
  across zones 603, 609
  cloud computing 605, 610
HackerOne 203
hackers 121. *See also* penetration testing
hacktivists 120, 122
hands-on activities 953
hard disks
  backups 326
  encryption 473
  forensic acquisition of 848
  hardening 157–159
  redundancy
    definition of 315–316
    multipath 319
    Redundant Array of Inexpensive Disks (RAID) 315–316
  self-encrypting 475–476
hardening
  applications 471
  hard disks 157–159
  operating systems 473–474
hardware root of trust 476–477
hardware security modules (HSMs) 478, 524, 587, 656
Hardware Shield 851
hashcat 749
HashCorp Nomad 240
Hashed Message Authentication Mode (HMAC) 295–296, 551–552

hashing 218–219, 463, 856–857
avalanche effect 463
collisions 463
definition of 398–399
Digital Signature Algorithm (DSA) 396
Elliptic Curve Digital Signature Algorithm (ECDSA) 551–552
Hashed Message Authentication Mode (HMAC) 295–296
Message Digest Algorithm 5 (MD5) 55, 219
padding 463
pass the hash 89–90
salting 47, 82
Secure Hash Algorithm (SHA) 55, 463, 551–552
SHA-256 463
HAVA (Help America Vote Act) 880
head command 733
headers, HTTP (Hypertext Transfer Protocol) 465–466
Health Insurance Portability and Accountability Act (HIPAA) 453, 880, 940, 944
heat maps 559
heating, ventilation, and air conditioning (HVAC) 352–353
Help America Vote Act (HAVA) 880
heuristic-based analysis 521
heuristic-based intrusion detection 521
HID Global 629
HIDSs (host intrusion detection systems) 215, 456, 578, 586
high availability (HA) 329–330
across zones 603, 609
cloud computing 605, 610
hijacking
BGP 535–536
blind 84
cookie 465
session 54, 83, 465, 601
TCP/IP 84
URL 44
hijacking, domain 108
HIPAA (Health Insurance Portability and Accountability Act) 453, 880, 940, 944
HIPSs (host intrusion prevention systems) 454–455, 523
Hitachi 476
HMAC (Hashed Message Authentication Mode) 295–296, 551–552
HMAC-based one-time password (HOTP) 295–296
HMI (human-machine interface) 341
hoaxes 19
holds, legal 842
HOME environment variable 740
homomorphic encryption 417
homomorphic steganography 417
honeyfiles 223
honeypots 221–223
hop-by-hop headers (HTTP) 466
horizontal privilege escalation 67–68
host command 716
host intrusion detection systems (HIDS) 215, 456, 578, 586
host intrusion prevention systems (HIPSs) 454–455, 523
host security. *See also* application security
antimalware 452
antivirus software 451
boot integrity
boot attestation 460–461
definition of 458–459
measured boot 459–460
Unified Extensible Firmware Interface (UEFI) 459
data loss prevention (DLP) 453
databases 461–462
endpoint 451, 452–453
hashing 463

host intrusion detection systems (HIDS) 215, 456, 578, 586
host intrusion prevention systems (HIPSs) 454–455, 523
host-based firewalls 457–458
next-generation firewall (NGFW) 453–454
salting 462–463
Host-based IPSs (HIPSs) 523
hot aisles 386
hot sites 221
hotfixes and patches 160–164, 179–180, 362, 474–475
HOTP (HMAC-based one-time password) 295–296
hotspots 585
hping command 717
Hping.org 717
HSMs (hardware security modules) 478, 524, 587, 656
HTTP (Hypertext Transfer Protocol) 465–466, 577
HTML5 505–508
HTTPS 82, 268, 436–437, 577
human resources (HR) personnel 901
human-machine interface (HMI) 341
HUMINT (human intelligence) 18
HVAC (heating, ventilation, and air conditioning) 352–353
hybrid attacks 749
hybrid cloud 140, 233
hyper-jacking 248
Hypertext Transfer Protocol. *See* HTTP (Hypertext Transfer Protocol)
hypervisors 325
attacks 601
hypervisor-based keyloggers 42

I

IA (information assurance). *See* risk management
IaaS (infrastructure as a service) 139, 231, 603, 853
IaC (infrastructure as code) 241–243, 260
IACS 342
IACS (industrial automation and control systems) 342, 343
IAM (identity and access management) 633
identity and access lifecycle 633–635
account audits 635
disablement 635
privileges provisioning 635
registration and identity validation 633–635
policy 605
IBM
AppScan 204
Data Encryption Standard (DES) 412
QRadar 526
IC (integrated circuit) cards 373
ICCIDs (unique serial numbers) 360
ICS-CERT (Industrial Control Systems Cyber Emergency Response Team) 362
ICSs (industrial control systems) 353–354
identification phase, incident response (IR) 763
identity. *See also* authentication; certificates; passwords
discovery of 623–624
federation 623
identity and access management (IAM) 633–635
identity and access lifecycle 633–635
policy 605
identity fraud 17, 638
baiting 19
credential harvesting 18
hoaxes 19
identity theft 940
impersonation/pretexting 19
invoice scams 17
reconnaissance 18
typo squatting 20, 44
watering hole attacks 20, 85

identity providers (IdPs) 292, 623–624, 661
Secure Shell (SSH) keys 628
smart cards 629
tokens 627–628
Identity Services Engine (ISE) 590
IdPs (identity providers) 292, 623–624, 661
IDSs (intrusion detection systems). *See* HIDSs (host intrusion detection systems); network intrusion detection systems (NIDSs)
IEEE 802.1X standard 510, 553–556, 562, 664–667, 673
IETF (Internet Engineering Task Force)
IPFIX (Internet Protocol Flow Information Export) 187
RFC (request for comments) 128
ifconfig command 710–711
IIS (Internet Information Services) 146, 697, 794
IKE (Internet Key Exchange)
IKEv1 Phase 1 negotiation 498–501
IKEv1 Phase 2 negotiation 501–503
IKEv2 504–505
image backups 326
image steganography 416–417
IMAP (Internet Message Access Protocol) 438–439
IMEI (international mobile equipment identity) 49, 584
immutability 263
impact assessment 184, 920, 921, 948
impersonation 19
implicit deny 528, 680
impossible travel time 639
IMSI (international mobile subscriber identity) encryption 358, 584
in-band SQL injection 73
incident response (IR) plans
business continuity plans (BCPs) 773–774, 929
communication plans 771–772
continuity of operations plans (COOPs) 774–775, 929
cyber kill chain 770–771
data retention policies 775–776
definition of 760–761
Diamond Model of Intrusion Analysis 768–770
disaster recovery plans (DRPs) 772–773
exercises
simulations 766–767
tabletop 765–766
walkthrough 766
incident response teams 175, 760, 775–776
MITRE ATT&CK framework 18, 128–129, 176, 205, 223, 767–768
process and lifecycle
containment 763–764
eradication 764
identification 763
lessons learned 764–765
overview of 761–762
preparation 762–763
recovery 764
stakeholder management 771–772
incident response (IR) teams 175, 760, 775–776
incremental backups 326, 328
indicators of compromise (IoCs) 123, 762, 832, 853
industrial automation and control systems (IACS) 342, 343
industrial camouflage 377
Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) 362
industrial control systems (ICSs) 341–343, 353–354
Industry 4.0 342
influence campaigns 21

information assurance (IA). *See* risk management
information lifecycle 947–948
Information Sharing and Analysis Centers (ISACs) 123–125
Information Society Directive 220
information systems security officers (ISSOs) 930, 947
Information Technology Infrastructure Library (ITIL) 882
information technology operations 263
InfraGard 128
infrastructure as a service (IaaS) 139, 231, 603, 853
infrastructure as code (IaC) 241–243, 260
inherent risk 921
inheritance, of permissions 644–646
Initial Contact 501
initialization vectors (IVs) 103, 403
injection 70
  code 149, 273–274, 276
  DLL (dynamic link library) 74
  LDAP (Lightweight Directory Access Protocol) 74, 144
  SQL (Structured Query Language) 54, 70–74, 273–274, 464
  XML (Extensible Markup Language) 74–75
inline prevention detection systems (IPSs) 523–524
input handling 79–82
input validation 80, 81, 267–268, 464
INSERT INTO statement 70
inspection, API 607, 610
instance awareness 608, 611
insurance, cybersecurity 918
integer overflows 77, 271
integrated circuit (IC) cards 373
integration
  API 607, 610
  cloud computing 604, 609
  continuous 279
integrity 289
  boot
    boot attestation 460–461
    definition of 458–459
    measured boot 459–460
    Unified Extensible Firmware Interface (UEFI) 459
  forensic acquisition 856
  integrity control 378
  measurement of 184, 261, 887
Intel Hardware Shield 851
intellectual property theft 917, 940
intelligence
  automated indicator sharing (AIS) 125
  Information Sharing and Analysis Centers (ISACs) 123–125
  intelligence fusion 177
  MITRE ATT&CK framework 18, 128–129, 176, 205, 223, 767–768
  research sources 127–128
  strategic 860
  Structured Threat Information eXpression (STIX) 125–127
  Trusted Automated eXchange of Indicator Information (TAXII) 125–127
  vulnerability databases 125
interconnection security agreements (ISAs) 903
intermediate certificate authorities 696
internal actors 122
internal information 905
internal risk 917. *See also* risk management
international mobile equipment identity (IMEI) 49, 584
international mobile subscriber identity (IMSI) 49, 358, 584
International Organization for Standardization (ISO) 881, 884, 893
Internet Engineering Task Force (IETF)

IPFIX (Internet Protocol Flow Information Export) 187
RFC (request for comments) 128
Internet Information Services (IIS) 146, 697, 794
Internet Key Exchange. *See* IKE (Internet Key Exchange)
Internet Message Access Protocol (IMAP) 438–439
Internet of Things (IoT) 38, 98, 113, 344–346, 358, 414
Internet Protocol. *See* IP (Internet Protocol)
Internet Protocol Flow Information Export (IPFIX) 187, 524, 811–813
Internet Security Association and Key Management Protocol (ISAKMP) 497
Internet service providers (ISPs) 808
interviews, forensic 846–847
Intigriti 203
intranets 492–493
intrusion detection systems. *See* host intrusion detection systems (HIDS)
intrusion detection systems (IDSs). *See* network intrusion detection systems (NIDSs)
intrusion phase, cyber kill chain 770
intrusion prevention systems. *See* host intrusion prevention systems (HIPSs)
intrusive scans 182
Investigate 509–510
investigations, data sources for
bandwidth monitors 804
Internet Protocol Flow Information Export (IPFIX) 811–813
log files 789
application 792–793
authentication 789–796
Call Manager 799–800
Domain Name System (DNS) 795–796
dump files 797
journalctl 802
network 790
NXLog 803–804
security 793
Session Initiation Protocol (SIP) 800
syslog/rsyslog/syslog-ng 800–801
system 791–792
Voice over Internet Protocol (VoIP) 799–800
web server 794
metadata
in email 808
in files 809
on mobile devices 808
types of 805–806
on web pages 808–809
NetFlow 809–810
protocol analyzers 813
Security Information and Event Management (SIEM)
alerts 788
correlation 788–789
dashboards 786–789
sensitivity 788
sensors 787
trends 788
sFlow 810–811
vulnerability scan output 785–786
invoice scams 17
IoCs (indicators of compromise) 123, 762, 832, 853
IoT (Internet of Things) 38, 98, 113, 344–346, 358, 414
IP (Internet Protocol). *See also* IPsec
addresses
IPv4 443–444
IPv6 536–537
virtual 488

configuration management 213
IP proxy 514
IP scanners
arp command 721–722
Cuckoo 731–732
curl command 724–725
definition of 721
dnsenum 728–729
Nessus 730–731
route command 723–724
scanless 727–728
sn1per 726–727
theHarvester 725–726
IP-Box 850–851
ipconfig command 710
IPFIX (Internet Protocol Flow Information Export) 187, 524, 811–813
IPsec 247, 437–438, 497. *See also* IKE (Internet Key Exchange)
attributes 501–502
Authentication Header (AH) 437
Encapsulating Security Payload (ESP) 437, 503
IKEv1 Phase 1 negotiation 498–501
IKEv1 Phase 2 negotiation 501–503
IKEv2 504–505
modes 438, 503
passthrough 501
IPSs (intrusion prevention systems). *See* HIPSs (host intrusion prevention systems); network intrusion detection systems (NIDSs)
IR. *See* incident response (IR) plans
iris recognition 301
ISACA COBIT framework 882
ISACs (Information Sharing and Analysis Centers) 123–125
ISAKMP (Internet Security Association and Key Management Protocol) 497
ISAs (interconnection security agreements) 903
ISE (Identity Services Engine) 590
ISO (International Organization for Standardization) 881, 884, 893
isolation 491, 562, 830
ISPs (internet service providers) 808
ISSOs (information systems security officers) 930, 947
issuers, certificate 692
IT contingency planning (ITCP) 929
IT security frameworks 881–884
ITIL (Information Technology Infrastructure Library) 882
ITU-T X.690 encoding formats 697
IVs (initialization vectors) 103, 403

J

jamming 102, 561–562
Japan's Personal Information Protection Act (JPIPA) 220
JavaScript Object Notation (JSON) injection 273–274
JavaScript-based keyloggers 43
job rotation 898, 900
John the Ripper 44, 749
journalctl 802
jump servers 514
jurisdictional forensic intervention 855

K

Kali forensics 850
Kali Linux 415, 953
Katacoda 239
KBA (knowledge-based authentication) 625, 656–657
KDC (key distribution center) 668
KE (Key Exchange) 500
Kerberoasting TGS 292
Kerberos 82–83, 89, 292, 553, 668–670, 673

kernel-based keyloggers 42
Key Exchange (KE) 500
.key file extension 697
key recovery agents 699
key signing keys (KSKs) 427
keyloggers 42–43, 108, 113
keys 688
  ephemeral 403
  escrow 699
  generation algorithms for 395
  key distribution center (KDC) 668
  key exchanges 399
  key signing keys (KSKs) 427
  length of 396
  mobile device management (MDM) 577–578
  password 655
  personal unblocking keys (PUKs) 360
  Public Key Cryptography Standards (PKCS) 412
  public/private 436–437
  Secure Shell (SSH) 625, 628
  stretching 397
  zone signing keys (ZSKs) 427
kiting, domain name 109–110
knowledge-based authentication (KBA) 625, 656–657
known environment/white box testing 198, 468–469
KSKs (key signing keys) 427
Kubernetes 239–240, 279–280

L

L0phtCrack 47
L2F (Layer 2 Forwarding Protocol) 508
L2TP (Layer 2 Tunneling Protocol) 494, 505–508
LANG environment variable 740
last known good configuration (LKGC) 329
lateral movement 201, 770
lateral traffic 492
laws 879–880
Layer 2 attacks
  ARP cache poisoning 105
  MAC cloning attacks 106
  MAC flooding attacks 106
  security best practices 106–107
Layer 2 Forwarding Protocol (L2F) 508
Layer 2 security 512
  Bridge Protocol Data Unit (BPDU) guard 512
  DHCP snooping 512–513
  loop protection 512
  MAC filtering 513
Layer 2 Tunneling Protocol (L2TP) 494, 505–508
LCP (Link Control Protocol) 44
LDAP (Lightweight Directory Access Protocol)
  injection attacks 144, 273–274, 291, 442, 667–670
  Lightweight Directory Access Protocol over SSL (LDAPS) 432
LDAPS (Lightweight Directory Access Protocol over SSL) 432
leaks, memory 78, 88
LEAP (Lightweight EAP) 666
least functionality 152
least privilege 264, 630, 681, 908
least significant bit (LSB) steganography 416–417
least-trusted zones 825
ledgers, public 409–410
legacy platforms 165
legal hold 842
lessons learned phase, incident response (IR) 764–765
libraries, third-party 265
licensing 918
lifecycle
  identity and access 633–635
    account audits 635
    disablement 635

privileges provisioning 635
registration and identity validation 633–635
incident response (IR)
containment 763–764
eradication 764
identification 763
lessons learned 764–765
overview of 761–762
preparation 762–763
recovery 764
information 947–948
penetration testing 199–202
lighting, security 380
lightweight cryptography 414–415
Lightweight Cryptography Project 415
Lightweight Directory Access Protocol. *See* LDAP (Lightweight Directory Access Protocol)
Lightweight Directory Access Protocol over SSL (LDAPS) 432
Lightweight EAP (LEAP) 666
Link Control Protocol (LCP) 44
Linux
Kali Linux 415
Linux Kernel 236
System Monitor 542
lists
allow 467, 578, 583, 822
block/deny 467–468, 822–823
certificate revocation 829
live boot media 329
live box forensics 858
load balancers 319–320
load balancing
active/active 488
active/passive 488
definition of 488
scheduling 488
Virtual IP address 488
Local Security Authority Subsystem Service (LSASS) 47–48
locality attribute (certificates) 692
Lockheed Martin 770
locks and lockout programs 378–379, 579, 639
log collectors 186
log files 789
aggregation 186
analytics 383
application 792–793
audit 869–870
authentication 789–796
Call Manager 799–800
collection of 186
correlation of 186
Domain Name System (DNS) 795–796
dump files 797
emergency preparedness 383
event 845–846
journalctl 802
network 790, 852–853
normalization of 186
NXLog 803–804
review 182
risk 920
security 383, 793
Session Initiation Protocol (SIP) 800
syslog/rsyslog/syslog-ng 800–801
system 791–792
visitor 383
Voice over Internet Protocol (VoIP) 799–800
web server 794
logger command 737–738
logic bombs 39–40
logistics, SCADA control systems 343
loop protection 512
LS_COLORS environment variable 740
LsaLogonUser 90
LSASS (Local Security Authority Subsystem Service) 47–48
LSB (least significant bit) steganography 416–417

M

MaaS (monitoring as a service) 139, 232
MAC (mandatory access control) 588, 676, 679, 905
MAC (media access control) 511
  addresses 511
  cloning attacks 106
  filtering 513
  flooding attacks 106
  spoofing 101
MACB (Modified, Accessed, Changed, and Birth) times 844
machine certificates 696
machine learning. *See* AI/ML (artificial intelligence and machine learning)
macOS Activity Monitor 542
macros 113
MACs (message authentication codes) 399, 410
MAIL environment variable 740
malicious software. *See* malware
Maltego 203
malware 113
  antimalware 452
  backdoors 42–43
  bots and botnets 37–38, 111–112
  cryptomalware 33–34
  definition of 33
  delivery mechanisms 43–45
  fileless viruses 37
  keyloggers 42–43
  logic bombs 39–40
  malvertising 40
  mobile device security countermeasures 580
  permanent damage from 45
  potentially unwanted programs (PUPs) 40–42
  ransomware 33–34, 111–112
  spyware 40–42
  time bombs 39
  Trojans 35, 104, 113
  worms 36–37
MAM (mobile application management) 585–587
managed detection and response (MDR) 234
managed power distribution units (PDUs) 322–323
managed security service providers (MSSPs) 233–234
managed service providers (MSPs) 233–234
management
  managerial controls 868
  roles and responsibilities 945–947
Management Information Bases (MIBs) 436
mandatory access control. *See* MAC (mandatory access control)
mandatory vacation policies 900
man-in-the-middle (MITM) attacks. *See* on-path (man-in-the-middle) attacks
manipulating files. *See* file manipulation
manual code review 470
manufacturing, SCADA control systems 342
mapping
  many-to-one 690
  one-to-one 690
masking, data 945
Mavituna Security Netsparker 204
maximum transmission unit (MTU) discovery 717
MBR (master boot record) 35–36, 851
MD5 algorithm 55, 219
MDM (mobile device management) 152, 574–576, 825–826, 908
  application and content management 576–578

bring-your-own-device (BYOD) 215, 572, 574–576, 581, 588–590, 826, 898
choose-your-own-device (CYOD) 588–590
corporate-owned, personally enabled (COPE) 572, 588–590
enforcement and monitoring 581–585
metadata 808
mobile application management (MAM) 585–587
SEAndroid 588
security concerns and countermeasures 578–581
unified endpoint management (UEM) 587–588
virtual desktop infrastructure (VDI) 589
MDR (managed detection and response) 234
mean time between failures (MTBF) 926
mean time to failure (MTTF) 926
mean time to repair (MTTR) 926
measured boot 459–460
Measurement System Analysis (MSA) 904
MEC (multi-access edge computing) 235
media access control. *See* MAC (media access control)
medical systems 347
MEIDs (mobile equipment identifiers) 49, 584
memdump 745
memorandum of understanding (MOU) 903
memory management 265. *See also* buffer overflows
ARP cache poisoning 105
content addressable 106
leaks 78, 88, 271
random-access memory (RAM) 849–850
runtime 477
static random-access memory (SRAM) 340
virtual 850
vulnerabilities 77–78, 149, 271–272, 275
memory-injection-based keyloggers 43
Men & Mice Logeater 796
Mentor Nucleus RTOS 347
message authentication codes (MACs) 399, 410
Message Digest Algorithm 5 (MD5) 55, 219
metadata
in email 808
in files 809
on mobile devices 808
types of 805–806
on web pages 808–809
Meterpreter scripts 90
MFA (multifactor authentication) 304–306, 579, 656–657
MFPs (multifunction printers) 354
MicroSD hardware security modules (HSMs) 587
microsegmentation 240–241, 489–490
microservices 236–240
Microsoft
Active Directory (AD) 291–292
Azure 232–233, 603, 853
Cluster Server 488
Defender Antivirus 823–824
Disk Defragmenter 158
Exchange 145
Forefront Identity Manager 658
Internet Information Services (IIS) 146
MS-CHAP 670–671
security advisories and bulletins 179
Security Bulletins 146
SQL Server 273
Visual Basic for Applications (VBA) 113
Web Application Proxy 516

Windows Defender Firewall 457
Windows Server 144
Mimikatz 90
minimal privilege 681
minimization, data 944–945
mirroring 316, 318, 537–538
mission-essential functions 929
mitigation 919, 921. *See also* segmentation
configuration changes 824
certificates, updating/revoking 829–830
content filter/URL filter 828–829
data loss prevention (DLP) 828
firewall rules 825
mobile device management (MDM) 825–826
containment 763–764, 830–831
endpoint security solutions 822
application approved lists 822
application block list/deny list 822–823
approved lists 822
block/deny lists 467–468, 578, 583, 822–823
quarantine 823–824
isolation 830
Security Orchestration, Automation, and Response (SOAR) 188–189, 832
playbooks 834
runbooks 833
MITRE Corporation 18, 458
ATT&CK framework 18, 128–129, 176, 205, 223, 767–768
Common Vulnerabilities and Exposures (CVE) 125, 146, 177
Common Weakness Enumeration 75
PRE-ATT&CK framework 18
MMS (Multimedia Messaging Service) 583, 585
mobile equipment identifiers (MEIDs) 49, 584
mobile solutions
Common Vulnerabilities and Exposures (CVEs) 571
connection methods and receivers 570
Bluetooth 570–571
cellular 572–573
Global Positioning System (GPS) 572, 584
near-field communication (NFC) 570–571
radio frequency identification (RFID) 571–572
satellite communications (SATCOM) 573
secure implementation best practices 573–574
mobile application management (MAM) 585–587
mobile device management (MDM) 215, 574–576
application and content management 576–578
bring-your-own-device (BYOD) 572, 574–576, 581, 588–590, 826, 898
choose-your-own-device (CYOD) 588–590
corporate-owned, personally enabled (COPE) 572, 588–590
enforcement and monitoring 581–585
mobile application management (MAM) 585–587
SEAndroid 588
security concerns and countermeasures 578–581
unified endpoint management (UEM) 587–588
virtual desktop infrastructure (VDI) 589
Modified, Accessed, Changed, and Birth (MACB) times 844

Modified Base Metrics 185
moisture detection systems 382
monitoring 537–538
  bandwidth 804
  continuous 139, 278
  file integrity monitors 542
  mobile device management (MDM) 581–585
  monitoring as a service (MaaS) 139, 232
  performance baselining 539–542
motion detection 382, 869–870
motion recognition 376
MOU (memorandum of understanding) 903
moves, MAC 511
MSA (Measurement System Analysis) 904
MS-CHAP 670–671
MSPs (managed service providers) 233–234
MSSPs (managed security service providers) 233–234
MTBF (mean time between failures) 926
MTTF (mean time to failure) 926
MTTR (mean time to repair) 926
MTU (maximum transmission unit) discovery 717
multi-access edge computing (MEC) 235
multicast addresses 537
multifactor authentication (MFA) 304–306, 579, 656–657
multifunction printers (MFPs) 354
multihomed connections 532
Multimedia Messaging Service (MMS) 583, 585
Multi-Party Coordination and Disclosure special interest group 180
multiparty risks 918
multipath I/O 319
multitenancy 601
Multi-User Multiple Input (MU-MIMO) 560–561
Mutiny Fuzzing Framework 269
mutual authentication 668–670
MySQL 273

N

NAC (network access control) 510–511, 871. *See also* 802.1X standard
name resolution 442–443
naming conventions 213
NarrowBand 358
NarrowBand-Internet of Things (NB-IoT) 358
NAS (network-attached storage) 326, 375
NAT (network address translation) 443–444, 501, 529, 562
Nation State attacks 346
National Cyber Awareness System (NCAS) 576
National Institute of Standards and Technology (NIST) 884
  cloud computing defined by 139
  Cybersecurity Framework (CSF) 884
  Digital Signature Algorithm (DSA) 396
  firewall guidelines 825
  isolation guidelines 830
  mobile device security guidelines 826
  National Vulnerability Database (NVD) 125, 177, 199
  NIST Cybersecurity Framework (CSF) 882
  Protecting Controlled Unclassified Information 828
  Risk Management Framework (RMF) 884
National Security Agency (NSA) 55, 498
National Vulnerability Database (NVD) 125, 177, 199
NAT-T (NAT Traversal) 501
NB-IoT (NarrowBand-Internet of Things) 358
NCAS (National Cyber Awareness System) 576

nCircle WebApp360 204
NDA (nondisclosure agreement) 901
near-field communication (NFC) 50, 100, 102–103, 570–571
negatives, false 181, 519, 520
Nessus 204, 730–731
net time command 669
netcat command 720–721
NetFlow 187, 525, 809–810
netstat command 668, 718–720
NetStumbler 99
network access control (NAC) 510–511, 871. *See also* 802.1X standard
network ACLs (access control lists) 535
network address translation (NAT) 443–444, 501, 529, 562
network and port scanners 182
network attached storage (NAS) 375
network attacks. *See also* network design, secure
  DDoS (distributed denial-of-service) 113
  DNS (Domain Name System)
    DDoS (distributed denial-of-service) 37–38, 54, 111–113, 601
    DNS amplification attack 112
    DNS poisoning 108–110
    domain hijacking 108
    domain name kiting 109–110
    domain reputation 110–111
    prevalence of 107
    URL redirection attacks 110
  Layer 2
    ARP cache poisoning 105
    MAC cloning attacks 106
    MAC flooding attacks 106
    security best practices 106–107
  malware 113
    backdoors 42–43
    bots and botnets 37–38, 111–112
    cryptomalware 33–34
    definition of 33
    delivery mechanisms 43–45
    fileless viruses 37
    keyloggers 42–43
    logic bombs 39–40
    permanent damage from 45
    potentially unwanted programs (PUPs) 40–42
    ransomware 33–34, 111–112
    spyware 40–42
    time bombs 39
    Trojans 35–36, 104, 113
    worms 36–37
  on-path attacks 54, 84–85, 103, 602
  password attacks
    brute-force 45
    dictionary-based 45
    password cracking 46
    password spraying 45
    plaintext/unencrypted 47–48
    rainbow tables 47
  replay attacks 82–85
  script execution 113
  wireless 98
    bluejacking 100
    bluesnarfing 99–100
    disassociation and deauthentication 101
    evil twin 98–99
    initialization vector (IV) 103
    jamming 102, 561–562
    near-field communication (NFC) 102–103
    radio frequency identification (RFID) 49, 102
    rogue access points 99
network controllers 144
network design, secure. *See also* firewalls; network attacks; network reconnaissance; network resilience
  access control lists (ACLs) 535, 643, 831
  broadcast storm prevention 512

Bridge Protocol Data Unit (BPDU) guard 512
DHCP snooping 512–513
loop protection 512
MAC filtering 513
DLP (data loss prevention) systems 215
Domain Name System (DNS) 509–510
load balancing
active/active 488
active/passive 488
definition of 488
scheduling 488
Virtual IP address 488
monitoring services 538–539
file integrity monitors 542
performance baselining 539–542
network access control (NAC) 510–511
network appliances 513–514
aggregators 526
hardware security modules (HSMs) 524
jump servers 514
network intrusion detection systems (NIDSs) 215, 223, 517–524, 870. *See also* network reconnaissance
network intrusion prevention systems (NIPSs) 99, 519, 869
network-based intrusion prevention system (NIPS) 518–524
proxy servers 514–516
sensors 524–525
network segmentation
application-based 489–490
east-west traffic 492
example of 489
extranets 492–493
intranets 492–493
microsegmentation 489–490
screened subnets 491
virtual local-area networks (VLANs) 490–491
zero trust 494
out-of-band management 510–511
port security 511, 537–538
route security 535–536
IPv6 536–537
port spanning/port mirroring 537–538
quality of service (QoS) 536
virtual private networks (VPNs) 507, 606
always-on VPN functionality 495
clientless versus client-based 497
concentrators 495
definition of 494
description of 494–496
example of 494–495
HTML5 508
IKEv1 Phase 1 negotiation 498–501
IKEv1 Phase 2 negotiation 501–503
IKEv2 504–505
IPsec 497
Layer 2 Tunneling Protocol (L2TP) 508
remote-access 496–497
site-to-site 495, 496–497
split tunneling 495–496
SSL (Secure Sockets Layer) 505–508
network forensic analysis tools (NFATs) 852–853
network interface card (NIC) teaming 320
network intrusion detection systems (NIDSs) 99, 215, 223, 517–518, 870. *See also* network reconnaissance
advantages/disadvantages 519–520
anomaly-based analysis 521–523
definition of 519–520
heuristic-based analysis 521
inline versus passive 523–524
promiscuous mode 517
signature-based 519–520

stateful pattern-matching recognition 521
network intrusion prevention systems (NIPSs) 99, 519, 869
network logs 790
Network Policy Server (NPS) 495
network reconnaissance 18, 770
active 204–205
definition of 707
dig 709–710
hping 717
ifconfig 710–711
ipconfig 710
netcat 720–721
netstat 718–720
nmap 711–714
nslookup 709–710
passive 203–204
pathping 716–717
ping 714–716
ping6 716
tracert/traceroute 707–709
network resilience
definition of 319
load balancers 319–320
network interface card (NIC) teaming 320
network segmentation. *See* segmentation
Network Time Protocol (NTP) 112, 440, 490, 790
Network Time Security (NTS) 440
network video recorders (NVRs) 375
network-attached storage (NAS) 326
network-based application layer firewalls 530
New Technology File System (NTFS) 156, 646, 850. *See also* permissions
Nexpose 204
next-generation firewall (NGFW) 453–454, 524
next-generation IPS systems (NGIPSs) 523
NFATs (network forensic analysis tools) 852–853
NFC (near-field communication) 50, 570–571
NFC (near-field communication) attacks 102–103
NGFW (next-generation firewall) 453–454, 524
nginx 236, 794
NGIPSs (next-generation IPS systems) 523
NIC (network interface card) teaming 320
NIDSs (network intrusion detection systems) 99, 215, 223, 517–518, 869, 870. *See also* network reconnaissance
advantages/disadvantages 519–520
anomaly-based analysis 521–523
definition of 519–520
heuristic-based analysis 521
inline versus passive 523–524
promiscuous mode 517
signature-based 519–520
stateful pattern-matching recognition 521
Nikto 204
Nimda 37
NIPSs (network intrusion prevention systems) 99, 523, 869
advantages/disadvantages 519–520
anomaly-based analysis 521–523
definition of 518–520
false positives/false negatives 519
heuristic-based analysis 521
inline versus passive 523–524
signature-based 520–521
NIST (National Institute of Standards and Technology) 396, 881
Nmap 204, 527, 711–714
noise detection 382
Nomad 240

nonces 82–83, 500
noncredentialed vulnerability scans 182
nondisclosure agreements (NDAs) 901
nonintrusive vulnerability scanners 182
non-persistence 328–329
nonrepudiation 859–860
normalization 186, 273–274
NoSQL databases 273–274
notifications
  of privacy and data breaches 941
  public 940
  push 299
Novec 1230 381
NPS (Network Policy Server) 495
NSA (National Security Agency) 55, 498
nslookup command 709–710
NT LAN Manager (NTLM) 89
NTFS (New Technology File System) 156, 646, 850. *See also* permissions
NTLM (NT LAN Manager) 89
NTP (Network Time Protocol) 112, 440, 490
NTS (Network Time Security) 440
Nucleus RTOS 347
null pointer dereferences 75, 271–272
NVD (National Vulnerability Database) 125, 177, 199
NVRs (network video recorders) 375
NXLog 803–804

O

Oakley 497
OAS (OpenAPI Specification) 87
OAuth 292, 578, 661–662
obfuscation 79, 265, 770
object detection 376
object identifiers (OIDs) 691
OCIL (Open Checklist Interactive Language) 885
OCSP (Online Certificate Status Protocol) 691, 698
OEM (original equipment manufacturer) 459
OFB (Output Feedback) mode 407
offboarding policies 575, 899, 900
Office of Personnel Management (OPM) attack 300–301
offline backups 326
offline password cracking 46
off-premises services 234
offsite storage 327
Off-The-Record Messaging 400–401
OIDC (OpenID Connect) 663–664
OIDs (object identifiers) 691
OLDPWD environment variable 740
onboarding policies 575, 899, 900
one-time passwords (OTPs) 627
  HMAC-based 295–296
  time-based 295
one-to-one mapping 690
one-way functions 219
online backups 326
Online Certificate Status Protocol (OCSP) 691, 698
online password cracking 46
on-path (man-in-the-middle) attacks 54, 84–85, 103, 602
on-premises environments, vulnerabilities in 137–143
on-premises services 234
Opal 476
Open Checklist Interactive Language (OCIL) 885
Open Network Environment 882
open permissions 150
open ports/services 471–472
Open Source Security Testing Methodology Manual (OSSTMM) 199
Open Systems Interconnection (OSI) model 103, 614, 615
Open vSwitch Database Management Protocol (OVSDB) 243

Open vSwitch (OVS) 243
Open Vulnerability and Assessment Language (OVAL) 164, 885
Open Web Application Security Project. *See* OWASP (Open Web Application Security Project)
Open1X 554
OpenAPI Specification (OAS) 87
OpenCv 416
OpenDNS 509–510
OpenFlow 243, 882
OpenID 663–664
open-source intelligence (OSINT) 7–8, 18, 120–121, 124, 203
OpenSSL 236, 741–742
OPENSSL_CONF environment variable 740
operating systems (OSs)
  forensic acquisition 850
  hardening 473–474
  trusted operating systems (TOSs) 905
operation, modes of (encryption)
  authenticated 404
  Cipher Block Chaining (CBC) 405
  Cipher Feedback (CFB) 406
  Counter (CTR) 404, 408–409
  Electronic Code Book (ECB) 404
  Output Feedback (OFB) 407
  unauthenticated 404
operational controls 868, 869
operational expenditure (OpEx) 598
operational technology (OT) 113
The Orange Book 674
order of volatility 848
organization attribute (certificates) 692
organizational incidents 775
organizational security. *See also* forensics, digital; incident response (IR) plans
  benchmarks and secure configuration guides 885–888
  data sanitization 748–749
  exploitation frameworks 747–748
  file manipulation 732–733
    cat command 734–735
    chmod command 736–737
    grep command 735–736
    head command 733
    logger command 737–738
    tail command 734
  IP scanners
    arp command 721–722
    Cuckoo 731–732
    curl command 724–725
    definition of 721
    dnsenum 728–729
    Nessus 730–731
    route command 723–724
    scanless 727–728
    sn1per 726–727
    theHarvester 725–726
  IT security frameworks 881–884
  network reconnaissance
    definition of 707
    dig 709–710
    hping 717
    ifconfig 710–711
    ipconfig 710
    netcat 720–721
    netstat 718–720
    nmap 711–714
    nslookup 709–710
    pathping 716–717
    ping 714–716
    ping6 716
    tracert/traceroute 707–709
  packet capture and replay
    definition of 742
    Tcpdump 742–743

Tcpreplay 742
Wireshark 743
password crackers 748–749
policies
acceptable use 898, 900
asset management 909–910
breadth and scope of 897
change management/change control 909
classification and governance 904–905
clean desk policy 23, 899, 900
credential 906–908
data retention 906
definition of 893
due care 900
due diligence 900
due process 900
job rotation 898, 900
mandatory vacations 898–899, 900
onboarding/offboarding 899, 900
privacy 897
procedures versus 893
separation of duties 898, 900
user education and awareness training 901–902
privacy and data breach consequences
data types and asset classification 941–942
fines 940
identity theft 940
impact assessment 948
information lifecycle 947–948
intellectual property theft 940
notifications 941
personally identifiable information (PII) 943
privacy enhancing technologies 944–945
privacy notices 949
protected health information (PHI) 944
reputation damage 940
security roles and responsibilities 945–947
terms of agreement 948
regulations and standards
company policies 878–879
General Data Protection Regulation (GDPR) 214, 220, 878–879, 947
laws 879–880
Payment Card Industry Data Security Standard (PCI DSS) 881
shell and script environments
definition of 738–740
OpenSSL 741–742
PowerShell 740
Python 741
Secure Shell (SSH) 739–740
organizational units (OUs) 692
organizational validation (OV) certificates 694
organized crime 120
original equipment manufacturer (OEM) 459
orthogonal frequency-division multiple access (OFDMA) 561
OSI (Open Systems Interconnection) model 103, 614, 615
OSINT (open-source intelligence) 7–8, 18, 120–121, 124, 203
OSSTMM (Open Source Security Testing Methodology Manual) 199
OT (operational technology) 113
OTA (over-the-air) technology 572–573, 583, 585
OTPs. *See* one-time passwords (OTPs)
out-of-band management 510–511
out-of-band SQL injection 73
Output Feedback (OFB) mode 407
outsourced code development 155
OV (organizational validation) certificates 694

OVAL (Open Vulnerability and Assessment Language) 164, 885
overflows
  buffer 75–76, 77, 149, 271–272, 275, 522
  integer 77, 271
over-the-air (OTA) technology 572–573, 583, 585
OVS (Open vSwitch) 243
OVSDB (Open vSwitch Database Management Protocol) 243
OWASP (Open Web Application Security Project) 204, 276–277
  OWASP Proactive Controls 276–277
  OWASP Testing Project 276
  OWASP Web Security Testing Guide 199
  top 10 vulnerabilities in web applications 70
  Top 10 Web Application Security Risks 277
  Zed Attack Proxy 204
owners, data 946
ownership, authentication by 625

P

P12/PFX format 697
PaaS (platform as a service) 139, 232, 853
PAC (proxy autoconfiguration) file 515
packet assemblers/disassemblers (PADs) 137–138
packet capture and replay 187
  definition of 742
  Tcpdump 742–743
  Tcpreplay 742
  Wireshark 743
packet filtering 528
packet sniffers 559
PacketFence 510
packet-switching exchanges (PSEs) 137–138
padding 463
PADs (packet assemblers/disassemblers) 137–138
Paessler PRTG 561–562
pagefiles, forensic acquisition of 849–850
palette modification 417
Palo Alto security advisories and bulletins 179
PAM (privileged access management) 678, 679
PAMs (pluggable authentication modules) 670
PAMs (Programmable Attribute Maps) 851
PANs (personal area networks) 570
PAP (Password Authentication Protocol) 670–671
parity, striping with (RAID) 316, 318
partially known environment 199
passive prevention detection systems (IPSs) 523–524
passive reconnaissance 203–204
pass-the-hash attacks 89–90
Password Authentication Protocol (PAP) 670–671
password crackers 748–749
passwords
  attacks
    brute-force 45
    dictionary-based 45
    password cracking 46
    password spraying 45
    plaintext/unencrypted 47–48
    rainbow tables 47
  cracking 748–749
  creating 636–638
  definition of 636
  HMAC-based one-time password (HOTP) 295–296
  mobile device management (MDM) 579, 582
  one-time passwords (OTPs) 627

Password Authentication Protocol (PAP) 670–671
password keys 655
password vaults 655
policies 906–907
system-generated 638
time-based one-time password (TOTP) 295
user-generated 638
Pastebin 18
patches and hotfixes 160–164, 179–180, 362, 474–475
PATH environment variable 740
pathping command 716–717
pattern-matching, stateful 521
payment methods, mobile 584, 586
PCI DSS (Payment Card Industry Data Security Standard) 453, 881
PDS (protective distribution system) 385
PDUs (power distribution units) 322–323
Peach 270
PEAP (Protected Extensible Authentication Protocol) 554, 556, 666
Pearson Test Prep practice test 954
peer to peer (P2P) networks 143
PEM (Privacy-enhanced Electronic Mail) 697
.pem file extension 697
penetration testing 121, 266
active reconnaissance 204–205
advantages of 197–198
bug bounties versus 202–203
cleanup 202
definition of 193, 197
exercise types 205–206
known environment 198
lifecycle 199–202
methodologies 199
partially known environment 199
passive reconnaissance 203–204
post-exploitation techniques 201
unknown environment 198–199
Penetration Testing Execution Standard (PTES) 199
Perfect Forward Secrecy (PFS) 399–400, 502
performance baselining 539–542
Performance Monitor tool 540–542
Performance tool 539
permissions 640–645
cloud computing 605, 610
inheritance 644–646
open 150
privilege creep 645
types of 646
persistence 201
personal area networks (PANs) 570
personal firewalls 534
personal identification numbers (PINs) 360, 579
Personal Identity Verification (PIV) cards 629
Personal Information Protection and Electronic Documents Act (PIPEDA) 220, 880
personal unblocking keys (PUKs) 360
personally identifiable information (PII) 82, 216–218, 268, 577, 856, 897, 901, 943
person-made disasters 924
personnel policies 377–378
acceptable use 898, 900
breadth and scope of 897
clean desk policy 23, 899, 900
data retention 906
definition of 893
due care 900
due diligence 900
due process 900
job rotation 898, 900
mandatory vacations 898–899, 900
onboarding/offboarding 575, 899, 900
personnel credential policy 906–908

privacy 897
procedures versus 893
separation of duties 898, 900
summary of 900
PFS (Perfect Forward Secrecy) 399–400, 502
pharming 14–15, 109
PHI (protected health information) 856, 944
phishing 9–12, 902
phone call authentication 299–300
physical controls 871–872
physical security 872
access control vestibules 372–373
air gap 384
alarms 374
attacks
card cloning 48–49
cloud-based attacks 52–55, 601–603
malicious flash drives 48
malicious USB cables 48
skimming 49–50
supply-chain attacks 51
badges 373, 382
bollards/barricades 370–371
cameras
centralized versus decentralized 375
closed-circuit television (CCTV) 376–377
motion recognition 376
object detection 376
drones 382–383
Faraday cages 383–384
fencing 380–381
fire suppression 381
industrial camouflage 377
lighting 380
locks 378–379
personnel 377–378
physical locks 379
protected cable distribution system 385
screened subnets 384
secure areas 385–386
secure data destruction 386–387
sensors 381–382
signage 374–375
USB data blockers 379–380
visitor logs 383
PIA (Privacy Impact Assessments) 948
piggybacking 15
PII (personally identifiable information) 82, 216–218, 268, 577, 856, 897, 901, 943
ping command 714–716
Ping of Death 88
ping6 command 716
PINs (personal identification numbers) 360, 579
PIPEDA (Personal Information Protection and Electronic Documents Act) 220, 880
PIR (Post Incident Review) 764–765
PIV (Personal Identity Verification) cards 629
pivoting 201
PKCS (Public Key Cryptography Standards) 412
PKI (public key infrastructure) 84–85, 556
certificate authorities (CAs) 556, 689–691, 829
certificates
attributes 691–692
chaining 696
expiration 693
formats 697
pinning 698
Subject Alternative Name 693
types of 694–696
definition of 685
key escrow 699
key management 688

key recovery agent 699
stapling 698
trust model 698
PKIX (Public Key Infrastructure Exchange) 694
plaintext 47–48
plans
business continuity 773–774, 929
communication 771–772
disaster recovery 772–773, 926
incident response (IR)
business continuity plans (BCPs) 773–774, 929
communication plans 771–772
continuity of operations planning (COOP) 774–775
cyber kill chain 770–771
data retention policies 775–776
definition of 760–761
Diamond Model of Intrusion Analysis 768–770
disaster recovery plans (DRPs) 772–773, 926
exercises 765–767
incident response teams 760, 775–776
MITRE ATT&CK framework 128–129, 176, 205, 223, 767–768
process and lifecycle 761–765
stakeholder management 771–772
platform as a service (PaaS) 139, 232, 853
platform configuration registers (PCRs) 294
playbooks 834
PLCs (programmable logic controllers) 341, 343
pluggable authentication modules (PAMs) 670
PlugX RAT 35
PMBOK (Project Management Body of Knowledge) 882
PNAC. *See* 802.1X standard
pointer dereferencing 75–76, 271–272
point-of-sale (POS) systems 353
Point-to-Point Tunneling Protocol (PPTP) 494, 558
poisoning
ARP (Address Resolution Protocol) 105, 722
DNS (Domain Name System) 108–110
policies 878–879
account 633
asset management 909–910
change management/change control 909
classification and governance 904–905
credential 906–908
data retention 775–776, 906
definition of 893
group policy objects (GPOs) 474
Identity and Access Management (IAM) 605
personnel
acceptable use 898, 900
breadth and scope of 897
clean desk policy 23, 899, 900
due care 900
due diligence 900
due process 900
job rotation 898, 900
mandatory vacations 898–899, 900
onboarding/offboarding 575, 899, 900
personnel credential policy 906–908
privacy 897
separation of duties 898, 900
summary of 900
procedures versus 893
resource 246, 603, 609
user education and awareness training 901–902
POP (Post Office Protocol) 438–439

port security 106, 511. *See also* 802.1X standard
  open ports 471–472
  port numbers 441
  port spanning/port mirroring 537–538
  port taps 538
  port-based network access control (PNAC) 553–554
  protocols associated with 152–154
  Switched Port Analyzer (SPAN) 537–538
  vulnerabilities 151
portals, captive 559
PortSwigger Burp Suite Professional 204
POS (point-of-sale) systems 353
positives, true/false 181–182, 518, 520
POST (power-on self-test) 851
Post Incident Review (PIR) 764–765
Post Office Protocol (POP) 438–439
post-exploitation techniques 201
post-quantum cryptography 402
potentially unwanted programs (PUPs) 40–42
power distribution units (PDUs) 322–323
power loss 925
power resilience
  definition of 320
  dual supply 321–322
  generators 321
  managed power distribution units (PDUs) 322–323
  uninterruptible power source (UPS) 320–321
power-on self-test (POST) 851
PowerShell 113, 630, 740
PPTP (Point-to-Point Tunneling Protocol) 494, 558
PRE-ATT&CK 18
predictive analysis 127
preferred roaming list (PRL) 572
PREMIS (Preservation Metadata Implementation Strategies) 805
preparation phase, incident response (IR) 762–763
prepending 17
preservation, forensic 858
Preservation Metadata Implementation Strategies (PREMIS) 805
preshared key (PSK) 103, 551, 557–558
pretexting 19
preventative controls 869, 872
principals 623
printenv command 739
Privacy Act of 1974 879, 897
privacy breaches 220. *See also* identity
  data types and asset classification 941–942
  fines 940
  identity theft 940
  impact assessment 948
  information lifecycle 947–948
  intellectual property theft 940
  notifications of 941
  personally identifiable information (PII) 943
  privacy enhancing technologies 944–945
  privacy notices 949
  privacy policies 897
  protected health information (PHI) 944
  reputation damage from 940
  security roles and responsibilities 945–947
  terms of agreement 948
privacy enhancing technologies 944–945
Privacy Impact Assessments (PIA) 948
Privacy-enhanced Electronic Mail (PEM) 697
private cloud 140, 232–233

Private information 942
private information sharing centers 124
private keys 436
private subnets 606, 610
privilege
  creep 645
  escalation 67–68, 201, 770, 941
  least 681
  minimal 681
  provisioning 635
privileged access management (PAM) 678, 679
Privileges Required (PR) metric 183
PRNG (pseudorandom number generator) 49–50, 102, 571–572
procedures, policies versus 879, 893
production 260
Programmable Attribute Maps (PAMs) 851
programmable logic controllers (PLCs) 341, 343
programming testing methods
  compile-time errors 266–267
  fuzz testing 269–270
  input validation 80, 267–268
  penetration testing 266
  runtime errors 266–267
  static and dynamic code analysis 269
  stress testing 80, 266
programming vulnerabilities. *See* vulnerabilities
Project Management Body of Knowledge (PMBOK) 882
promiscuous mode 517
promiscuous ports 491
Proprietary information 942
protected cable distribution system 385
Protected Extensible Authentication Protocol (PEAP) 554, 556, 666
protected health information (PHI) 856, 944
protective distribution system (PDS) 385
protocol analyzers 813
protocols. *See individual protocols*
provenance, forensic 857–858
provisioning, application 260
proximity readers 373, 382
proxy autoconfiguration (PAC) file 515
proxy servers 514–516
  forward proxy 516
  reverse proxy 506–507, 516
  transparent proxy 516
PSEs (packet-switching exchanges) 137–138
pseudo-anonymization 945
pseudocodes 79
pseudorandom number generator (PRNG) 49–50, 102, 571–572
PSK (preshared key) 103, 551, 557–558
PTES (Penetration Testing Execution Standard) 199
public cloud 140, 232
public incidents 775
public information 905
public information sharing centers 124
public key algorithms 411
Public Key Cryptography Standards (PKCS) 412
public key infrastructure. *See* PKI (public key infrastructure)
Public Key Infrastructure Exchange (PKIX) 694
public keys 437
public ledgers 409–410
public notifications and disclosures 941
public subnets 606, 610
PUKs (personal unblocking keys) 360
pulping 386
pulverizing 387
PUPs (potentially unwanted programs) 40–42
purple team 205–206
push notifications 299
PWD environment variable 740
Python 113, 741

Q

QKD (quantum key distribution) 401–402
QoS (quality of service) 536
QRadar 526
qualitative risk management 921–922, 923
qualitative-to-quantitative score mapping 186
quality assurance (QA) 260, 261
quality of service (QoS) 536
Qualys 204
quantitative risk management 922–923
quantum cryptography 401–402
  communications 401–402
  computing 402
  definition of 401
quantum key distribution (QKD) 401–402
quarantine 823–824
quick mode, IKE 501

R

race conditions 79
Radamsa 269
radio, baseband 359
radio frequency identification (RFID) attacks 49, 102, 571–572
radio frequency interference (RFI) 383–384
RADIUS (Remote Authentication Dial-In User Service) 556–557, 672–673
RAID (Redundant Array of Inexpensive Disks) 315–316, 869
Rainbow Series 674
rainbow tables 47
RainbowCrack 47
RAM (random-access memory), forensic acquisition of 848–849
ransomware 33–34, 111–112
rapid application development (RAD) 262
Rapid STP 512
Rapid7 Nexpose 204
RAs (registration authorities) 690
RAS (Remote Access Service) 670–672
Raspberry Pi 339
RATs (remote access Trojans) 148
RBAC (role-based access control) 677, 679, 899
RC4 (Rivest Cipher 4) 412
RCE (remote code execution) 78, 146, 149, 275
RCS (Rich Communication Services) 585
RCSA (risk control self-assessment) 920
RDBMS (relational database management system) 273
RDP (Remote Desktop Protocol) 472
readers, proximity 373, 382
Real-Time Monitoring Tool (RTMT) 799
real-time operating systems (RTOSs) 347, 355
Real-Time Transport Protocol (RTP) 152. *See also* Secure Real-Time Transport Protocol (SRTP)
reception desks 378
recertification, user access 645
reconnaissance. *See* network reconnaissance
Recon-ng 203
recovery 764, 859
  disaster recovery planning 928–930
  recovery point objective (RPO) 929
  recovery time objective (RTO) 929
  restoration order 330–331
Red Hat security advisories and bulletins 179
red teams 205, 902
redaction 945
redirection attacks, URL 110
reduced sign-on 656
redundancy 926–927
  definition of 315

disk
definition of 315–316
multipath 319
Redundant Array of Inexpensive Disks (RAID) 315–316
diversity of 331–332
geographic dispersal 315
network
definition of 319
load balancers 319–320
network interface card (NIC) teaming 320
power
definition of 320
dual supply 321–322
generators 321
managed power distribution units (PDUs) 322–323
uninterruptible power source (UPS) 320–321
Redundant Array of Inexpensive Disks (RAID) 315–316, 869
refactoring, driver 89
reference architecture 884
Reflected XSS attacks 68
reflection 112
regedit command 472
registers, risk 920
registration, identity 633–635
registration authorities (RAs) 690
registry 472
regulations and standards
company policies 878–879
General Data Protection Regulation (GDPR) 214, 220, 878–879, 947
laws 879–880
Payment Card Industry Data Security Standard (PCI DSS) 881
regulatory forensic intervention 855
relational database management system (RDBMS) 273
Reliable Event Logging Protocol (RELP) 800
relying parties (SAML) 659
Remediation Level (RL) metric 185
remote access 442
Remote Access Service (RAS) 670–672
remote access Trojans (RATs) 148
remote-access VPNs 496–497
remote authentication
Challenge-Handshake Authentication Protocol (CHAP) 670–672, 673
RADIUS 556–557, 672–673
Remote Access Service (RAS) 670–672
TACACS+ 672–673
Remote Authentication Dial-In User Service (RADIUS) 556–557, 672–673
remote code execution (RCE) 78, 146, 149, 275
Remote Desktop Connection 152
Remote Desktop Protocol (RDP) 472
remote terminal units (RTUs) 341
remote wipe 579, 582
remotely operated underwater vehicles (ROVs) 353–354
removable media 123
replay, packet
definition of 742
replay attacks 82–85
Tcpdump 742–743
Tcpreplay 742
Wireshark 743
replication
cloud computing 605, 610
storage area networks (SANs) 323
virtual machines (VMs) 324–325
escape attacks 248–249
sprawl avoidance 247–248
Report Confidence (RC) metric 185
reports
after action report (AAR) 928–929
baseline 539–542

forensic 846
SIEM (Security Information and Event Management) 187
repositories, file/code 127
Representational State Transfer (REST) 86
reputation 110–111, 940
request for comments (RFC) 128
request forgeries 85–86
residual risk 919, 921
resilience 221–222
backups
cloud 326
comparison of 326–327
copy 326
differential 326, 328
disk 326
full 326, 328–331
image 326
incremental 326, 328
NAS (network-attached storage) 326
offsite storage 327
online versus offline 326
snapshot 326
tape 326
definition of 311
high availability (HA) 329–330
network
definition of 319
load balancers 319–320
network interface card (NIC) teaming 320
non-persistence 328–329
power
definition of 320
dual supply 321–322
generators 321
managed power distribution units (PDUs) 322–323
uninterruptible power source (UPS) 320–321
on-premises versus cloud 325
redundancy
definition of 315
disk 315–319
diversity of 331–332
geographic dispersal 315
network 319–320
power 320–323
Redundant Array of Inexpensive Disks (RAID) 315–316, 869
replication
storage area networks (SANs) 323
virtual machines (VMs) 247–249, 324–325
restoration order 330–331
scalability 279–280, 328
resolution, domain name 442–443
resource allocation, dynamic 607–608, 611
resource exhaustion 87–88
resource policies 246, 603, 609
resource records (RRs) 795
response and recovery controls 220–221
REST (Representational State Transfer) 86
RESTful APIs 240
restoration 158, 330–331
retention, risk 919
retention policies 775–776, 906
retina scanning 301
Retina Web Security Scanner 204
reuse, code 270
reverse proxy 506–507, 516
revert to known state 329
review, exam 953–954
review logs 182
reviews, configuration 182
revoking certificates 829
RFC (request for comments) 128
RFI (radio frequency interference) 383–384

RFID (radio frequency identification) attacks 49, 102, 571–572
Rich Communication Services (RCS) 583, 585
riding, session 602
rights management 219–220, 640–645
right-to-audit clauses 854
Rijndael. *See* Advanced Encryption Standard (AES)
risk management 155, 913
  business impact analysis 926–927
  disaster analysis 924–925
  disaster recovery planning 928–930
  external versus internal risk 917
  residual risk 919
  risk assessment 919–921
    control risk 921
    inherent risk 921
    qualitative 921–922, 923
    quantitative 922–923
    residual risk 921
    risk appetite 921
    risk awareness 921
    risk control assessment 920
    risk control self-assessment (RCSA) 920
    risk matrix/heat map 920
    risk mitigation 921
    steps of 919–920
  risk avoidance 918
  Risk Management Framework (RMF) 884
  risk matrix/heat map 920
  risk mitigation 919
  risk registers 920
  risk transference 918
  risk types 917–918
  strategies for 918–919
  supply chain risk management (SCRM) 920
  third-party risks 155–160
risky login 639
RMF (Risk Management Framework) 884
robot sentries 378
rogue access points 99
role-based access control (RBAC) 677, 679, 899
role-based training 902
roles and responsibilities, security 945–947
rolling codes 102
root accounts 150, 908
root certificate authorities 696
root certificates 696
root of trust 476–477
route command 723–724
route security 443, 535–536
  IPv6 536–537
  port spanning/port mirroring 537–538
  quality of service (QoS) 536
Routing and Remote Access Service (RRAS) 495
ROVs (remotely operated underwater vehicles) 353–354
RPO (recover point objective) 929
RRAS (Routing and Remote Access Service) 495
RRs (resource records) 795
RSA 412
rsyslog 800–801
RTMT (Real-Time Monitoring Tool) 799
RTO (recovery time objective) 929
RTOSs (real-time operating systems) 347, 355
RTP (Real-Time Transport Protocol) 152. *See also* Secure Real-Time Transport Protocol (SRTP)
RTUs (remote terminal units) 341, 343
rule-based access control 677, 678, 679
runbooks 833
runtime errors 81–82, 266–267
runtime memory 477

S

SaaS (software as a service) 138, 231, 444, 853
SAE (Simultaneous Authentication of Equals) 101, 551, 552
safes 385
salting 47, 82, 397–398, 462–463
SAM (Security Accounts Manager) 89
SAML (Security Assertion Markup Language) 659–661
Samsung 476
SAN (Subject Alternative Name) field 694–695
sandboxing 266, 452, 478–479
sanitization, data 748–749
sanitizing mobile devices 579
SANs (storage-area networks) 142, 323
Santos, Omar 953
Sarbanes–Oxley (SOX) 880, 882
SASE (Secure Access Service Edge) 582
SAST (static application security testing) 468–469
SATCOM (satellite communications) 573
SCADA (supervisory control and data acquisition) systems 341–343
scalability 279–280, 328
scanless 727–728
scans
  biometric. *See* biometric systems
  IP scanners
    arp command 721–722
    Cuckoo 731–732
    curl command 724–725
    definition of 721
    dnsenum 728–729
    Nessus 730–731
    route command 723–724
    scanless 727–728
    sn1per 726–727
    theHarvester 725–726
  vulnerability 785–786
    Common Vulnerability Scoring System (CVSS) 182–186
    false negative 181
    false positives 181
    how it works 180–181
    intrusive versus nonintrusive 182
    noncredentialed 182
SCAP (Security Content Automation Protocol) 883, 885–888
scheduling algorithms 488
SCP (secure copy) 456
screen locks 579
screened subnets 384, 491
script environments 278–279
  definition of 738–740
  OpenSSL 741–742
  PowerShell 740
  Python 741
  Secure Shell (SSH) 739–740
script kiddies 120
SCRM (supply chain risk management) 166, 920
Scrum 258
SDLC (software development lifecycle) 78, 261–262, 263–265, 468, 868
SDN (software-defined networking) 241–243, 882
SDV (software-defined visibility) 243
SD-WAN (software-defined wide-area network) 246
Seagate Technology 476
SEAndroid 588
search engine optimization (SEO) 808
SEC (Securities and Exchange Commission) 941
SECaaS (security as a service) 139
secrecy, forward 400–401
Secret information 905, 941–942
Secret Manager 604
secrets management 604, 609
Secure Access Service Edge (SASE) 582
secure areas 385–386

secure copy (SCP) 456
Secure File Transfer Protocol (SFTP) 434, 441
Secure Hash Algorithm (SHA) 55, 463, 551–552
Secure Key Exchange Mechanism (SKEME) 497
secure protocols. *See also individual protocols*
  definition of 426
  use cases
    directory services 442
    domain name resolution 442–443
    email and web 440
    file transfer 441
    network address allocation 443–444
    remote access 442
    routing and switching 443
    subscription services 444
    time synchronization 440
    voice and video 440
Secure Real-Time Transport Protocol (SRTP) 152, 430–431
Secure Shell (SSH) 427–428, 625, 628, 739–740
Secure Sockets Layer (SSL) 82–83, 436, 441
  certificate types 694–696
  SSL-based VPNs 505–508
  Transport Layer Security Inspection (TLSI) 215–216
Secure Web Gateway (SWG) 613, 614
Secure/Multipurpose Internet Mail Extensions (S/MIME) 428–429
Securities and Exchange Commission (SEC) 941
Security Accounts Manager (SAM) 89
security administrators 947
security as a service (SECaaS) 139
Security Assertion Markup Language (SAML) 292, 659–661
security assessments. *See also* SIEM (Security Information and Event Management)
  in cloud 598
    attacks 601–603
    threats 598–600
  risk 919–921
    control risk 921
    inherent risk 921
    qualitative 921–922, 923
    quantitative 922–923
    residual risk 921
    risk appetite 921
    risk awareness 921
    risk control assessment 920
    risk control self-assessment (RCSA) 920
    risk matrix/heat map 920
    risk mitigation 921
    steps of 919–920
  security advisories and bulletins 177–180
  Security Orchestration, Automation, and Response (SOAR) 188–189, 832
  threat hunting 175–180
  vulnerability scans
    credentialed versus noncredentialed 182
    false negatives 181
    false positives 181
    how it works 180–181
    intrusive versus nonintrusive 182
Security Content Automation Protocol (SCAP) 883, 885–888
security controls
  cloud
    API inspection and integration 607, 610
    compute 611
    high availability across zones 603, 609

integration and auditing 604, 609
network 606–607, 610
resource policies 603, 609
secrets management 604, 609
storage 605, 610
cloud computing
compute 607
container security 608–609
dynamic resource allocation 607–608, 611
instance awareness 608, 611
native versus third-party 615
security groups 607, 611
security solutions 611–614
summary of 608–609
virtual private cloud endpoint 608, 611
security incident response simulations (SIRS) 766–767
security incident response team (SIRT). *See* incident response (IR) teams
Security Information and Event Management. *See* SIEM (Security Information and Event Management)
security logs 383, 793
security officers 947
Security Onion 953
security operations centers (SOCs) 123, 175–176, 223, 379, 760, 762, 776
Security Orchestration, Automation, and Response (SOAR) 188–189, 832
playbooks 834
runbooks 833
security posture assessments (SPAs) 539
Security Requirements metrics 185
Security-Enhanced Linux (SELinux) 588, 676
SEDs (self-encrypting drives) 475–476
segmentation 607, 610, 831–832
application-based 489–490
in cloud 613, 615
east-west traffic 492
example of 489
extranets 492–493
intranets 492–493
microsegmentation 489–490
screened subnets 491
virtual local-area networks (VLANs) 490–491
zero trust 494
Segmented Integer Counter Mode (SRTP) 430
SEH (structured exception handling) 81, 267
SELECT statement 70
self-encrypting drives (SEDs) 475–476
self-signed certificates 695, 698
SELiux (Security-Enhanced Linux) 588
semi-authorized hackers 121
semicolon (;) 73
Sender Policy Framework (SPF) 110, 426
sensitive data exposure 82
Sensitive information 942
sensors 345, 381–382, 524–525, 787
sentiment analysis 188
Sentinel 204
SEO (search engine optimization) 808
separation of duties 898, 900
serial numbers, certificate 692
serverless architecture 243–244
servers 144
authentication 665
command-and-control [C2] 108
email 145
file 144
FTP 147–148
hardening 159–160
jump 514
Microsoft Cluster Server 488
network controllers 144
Network Time Protocol (NTP) 490
proxy 514–516
forward proxy 516

reverse proxy 506–507, 516
transparent proxy 516
virtual network computing (VNC) 632
web
log files 794
vulnerabilities 146–147
server-side execution 267
server-side request forgery (SSRF) 85–86
server-side validation 268
service accounts 629, 908
service nxlog start command 803
service providers (SPs) 292, 623, 661
service set identifiers (SSIDs) 98, 205, 532
service-level agreements (SLAs) 53, 273–274, 600, 902–903
services, open 471–472
services integration 246
session hijacking 54, 83, 465, 601
Session Initiation Protocol (SIP) 351, 431, 800
session replay 83
session riding 54, 602
session theft 83
SET (Social Engineering Toolkit) 10
SFC (System File Checker) command 158
sFlow 810–811
SFTP (Secure File Transfer Protocol) 434, 441
SHA (Secure Hash Algorithm) 55, 551–552
shadow IT 121
share permissions 646. *See also* permissions
shared accounts 629
shell and script environments
definition of 738–740
OpenSSL 741–742
PowerShell 740
Python 741
Secure Shell (SSH) 739–740
SHELL environment variable 740
shielding, application 471
shimming, driver 89
Shodan 203–204
Shor's algorithm 402
Short Message Service (SMS) 12, 296–297, 583, 585
shoulder surfing 14
shredding 386
side-channel attacks 54, 602
sideloading 581
SIEM (Security Information and Event Management) 186–188, 526, 869–870
alerts 788
correlation 788–789
dashboards 786–789
sensitivity 788
sensors 787
trends 788
SIFT workstation 850
signage 374–375
signatures, digital 395–396, 466–467, 520
signature verifying algorithms 395
signature-based intrusion detection 519–520
signing algorithms 395
SIM (subscriber identity module) cards 49, 360, 580, 584
Simple Network Management Protocol version 3 (SNMPv3) 434–436, 443
Simple Object Access Protocol (SOAP) 86
simulations 766–767
Simultaneous Authentication of Equals (SAE) 101, 551, 552
single loss expectancy (SLE) 922
single point of failure 156, 926
single quotation mark (') 73
single sign-on (SSO) 292, 373, 624, 658–659
sinkholes, DNS 223

SIP (Session Initiation Protocol) 351, 431, 800
SIRS (security incident response simulations) 766–767
SIRT. *See* incident response (IR) teams
site resiliency 221–222
site surveys 559, 561–562
sites, physical 385
site-to-site configuration 495
site-to-site VPNs 496–497
SKEME (Secure Key Exchange Mechanism) 497
SKEYID 500
skimming 49–50
SLAs (service-level agreements) 53, 273–274, 600, 902–903
SLE (single loss expectancy) 922
Sleuth Kit 850
smart cards 299–300, 625, 629
smart devices 345
smart factories 342
smart meters 350
S/MIME (Secure/Multipurpose Internet Mail Extensions) 428–429
smishing 12
SMS (Short Message Service) 12, 296–297, 583, 585
sn1per 726–727
snapshots 326, 851–852
SNMPv3 (Simple Network Management Protocol version 3) 434–436, 443
snmpwalk v3 command 436
snooping, DHCP 512–513
SOAP (Simple Object Access Protocol) 86
SOAR. *See* Security Orchestration, Automation, and Response (SOAR)
SOC (System and Organization Controls) 884
SoC (system on a chip) 356–357, 477, 571
social engineering attacks
  description of 7–9
  dumpster diving 13
  eliciting information 15–16
  hybrid warfare 22
  identity fraud 17
    baiting 19
    credential harvesting 18
    hoaxes 19
    impersonation/pretexting 19
    invoice scams 17
    reconnaissance 18
    typo squatting 20, 44
    watering hole attacks 20, 85
  influence campaigns 21
  pharming 14–15
  phishing and spear phishing 9–12
  piggybacking 15
  prepending 17
  principles of 21
  reasons for effectiveness 21
  shoulder surfing 14
  smishing 12
  Spam 13
  Spam over Internet Messaging (SPIM) 13
  tailgating 15
  user security awareness education 22–24
  vishing 12–13
  war-dialing 13
  whaling 9, 16–17
Social Engineering Toolkit (SET) 10
social media
  attacks and vulnerabilities 22, 123, 143
  as research source 128
social media analysis 899
SOCs (security operations centers) 123, 175–176, 223, 379, 760, 762, 776
software application development. *See* application development
software as a service (SaaS) 138, 231, 444, 853

software compliance/licensing 918
software development environments 257–260
software development lifecycle (SDLC) 78, 261–262, 263–265, 468, 868
software diversity 278
software integrity measurement 261
Software of Unknown Providence (SOUP) 347
software-defined networking (SDN) 241–243, 882
software-defined visibility (SDV) 243
software-defined wide-area network (SD-WAN) 246
SolarWinds 721, 789
solid-state drives (SSDs), forensic acquisition of 848
SOUP (Software of Unknown Providence) 347
sovereignty, data 214–215
SOX (Sarbanes–Oxley) 880, 882
Spam 13
Spam over Internet Messaging (SPIM) 13
SpamCop 13
SPAN (Switched Port Analyzer) ports 537–538
spanning, port 537–538
Spanning Tree Protocol (STP) 105, 512
spanning-tree portfast bpduguard command 512
SPAs (security posture assessments) 539
specialized embedded systems 346–347
  aircraft 348–350
  communication considerations
    5G 357–358
    baseband radio 359
    NarrowBand 358
    subscriber identity module (SIM) cards 360
    Zigbee 360–361
  constraints 361
    authentication 363
    compute 361–362
    cost 363
    crypto 362
    implied trust 363
    inability to patch 362
    network 362
    power 361
    range 363
  drones 353–354
  heating, ventilation, and air conditioning (HVAC) 352–353
  medical systems 347
  multifunction printers (MFPs) 354
  real-time operating systems (RTOSs) 355
  smart meters 350
  surveillance systems 355–356
  system on a chip (SoC) 356–357
  vehicles 347–348
  Voice over Internet Protocol (VoIP) 350, 799–800
speech recognition 302
SPF (Sender Policy Framework) 110, 426
SPI (stateful packet inspection) 528, 562
SpiderFoot 203
SPIM (Spam over Internet Messaging) 13
split tunneling 495–496
Splunk 526
spoofing
  ARP (Address Resolution Protocol) 105, 513
  MAC (media access control) 101, 106
sprawl avoidance 247–248
spraying, password 45
SPs (service providers) 292, 623, 661
spyware 40–42
SQL (Structured Query Language) 273
  SQL injection (SQLi) 54, 70–74, 273–274, 464, 602
  SQL Server 273
SquidProxies 514

SRAM (static random-access memory) 340
SRTP (Secure Real-Time Transport Protocol) 152, 430–431
SSAE (Statement on Standards for Attestation Engagements) 881, 883, 884
SSDs (solid-state drives), forensic acquisition of 848
SSH (Secure Shell) 427–428, 625, 628, 739–740
ssh command 427
SSIDs (service set identifiers) 98, 205, 532
SSL (Secure Sockets Layer) 82–83, 436, 441
  certificate types 694–696
  SSL-based VPNs 505–508
  stripping 88–89
  Transport Layer Security Inspection (TLSI) 215–216
SSL Inspection (SSSI) 215
SSO (single sign-on) 292, 373, 624, 658–659
SSRF (server-side request forgery) 85–86
SSSI (SSL Inspection) 215
staging 259
stakeholder management 771–772
standard load 540
standards. *See* regulations and standards
stapling 698
starvation attack, DHCP 513
state actors 120–121
state laws 879–880
stateful packet inspection (SPI) 528, 562
stateful pattern-matching recognition 521
stateless packet inspection 528
Statement on Standards for Attestation Engagements (SSAE) 881, 883, 884
statements, SQL (Structured Query Language) 70
static application security testing (SAST) 468–469
static code analysis 269, 468–469
static codes 298
static random-access memory (SRAM) 340
Stegais 415
steganography 415
  audio 415–416
  homomorphic 417
  image 416–417
  video 416
Steghide 415
stego-files 416
stewards, data 946
sticky sessions 489
STIX (Structured Threat Information eXpression) 125–127
storage
  cloud 610
    encryption 605
    high availability 605
    permissions 605
    replication 605
  secure 477
  storage DLP systems 215
  vulnerabilities 156
storage-area networks (SANs) 142, 323
Stored (persistent) XSS attacks 68
stored procedures 273
STP (Spanning Tree Protocol) 105, 512
strategic intelligence 860
stream ciphers 410
stress testing 80, 266
stretching, key 397
striping (RAID) 316, 317–318
  with dual parity 316, 318
  with parity 316, 318
  stripe and mirror 316, 319
stripping, SSL 88–89
structured exception handling (SEH) 81, 267

Structured Query Language. *See* SQL (Structured Query Language)
Structured Threat Information eXpression (STIX) 125–127
Stuxnet 363
Subject Alternative Name (SAN) 693, 694–695
subnets
  public/private 606, 610
  screened 384, 491
subscriber identity module (SIM) cards 49, 360, 580, 584
substitution 216, 416–417
supervisory control and data acquisition (SCADA) systems 341–343
supplicants 555, 665
supply chains
  attacks 51, 123, 156
  business partnership agreements (BPAs) 903
  supply chain risk management (SCRM) 166, 920
surge protectors 159
surveillance systems 355–356
surveys, site 559, 561–562
svStrike 850–851
Swagger (OpenAPI) 87
swap files, forensic acquisition of 849–850
SWG (Secure Web Gateway) 613, 614
Switched Port Analyzer (SPAN) ports 537–538
switching 443
symmetric encryption 411–413
synchronization 82–83
  email and web 440
  time 440
synchronization (SYN) packets 84
syslog 800–801
syslog-ng 800–801
System and Organization Controls (SOC) 884
System Information 161
system integration 155
system logs 791–792
System Monitor 542
system on a chip (SoC) 356–357, 477, 571
system owners 946–947
System Restore 158
systemd 802
system-generated passwords 638
systeminfo command 161

T

tables, rainbow 47
tabletop exercises 765–766
TACACS+ (Terminal Access Controller Access Control System Plus) 672–673
tactics, techniques, and procedures (TTPs) 128, 176, 767, 809
tags, evidence 845–846
tail command 734, 795
tailgating 15
Talos 347
tamper resistance 477
tape backups 326
taps, port 538
TAXII (Trusted Automated eXchange of Indicator Information) 125–127
TCB (trusted computing base) 676
TCG (Trusted Computing Group), Opal 476
Tcl 241
TCP (Transmission Control Protocol) 503
Tcpdump 742–743
TCP/IP hijacking 84
Tcpreplay 742
TCSEC (Trusted Computer System Evaluation Criteria) 674
teaming, network interface card (NIC) 320
teams, incident response (IR) 760, 775–776

Teardrop 88
technical controls 868, 869
Technical Guide to Information Security Testing and Assessment (NIST) 199
TEE (trusted execution environment) 476
telemetry, fake 223
temperature sensors 382
temporal groups 182
Temporal Key Integrity Protocol (TKIP) 552
temporary files 157
Tenable Network Security Nessus 204
TERM environment variable 740
Terminal Access Controller Access Control System Plus (TACACS+) 672–673
terms of agreement 948
testing 259
  black-box 80
  compile-time errors 266–267
  fuzz 80, 269–270
  gray-box 80
  input validation 80, 267–268
  known environment/white box 468–469
  penetration 121, 266
    active reconnaissance 204–205
    advantages of 197–198
    bug bounties versus 202–203
    cleanup 202
    definition of 193, 197
    exercise types 205–206
    known environment 198
    lifecycle 199–202
    methodologies 199
    partially known environment 199
    passive reconnaissance 203–204
    post-exploitation techniques 201
    rules of engagement 200
    unknown environment 198–199
  runtime errors 266–267
  static and dynamic code analysis 269
  stress 80, 266
  white-box 80
tethering 584
TGTs (ticket-granting tickets) 668
THC Hydra 749
theft
  disaster analysis 925
  identity 940
  intellectual property 917
  mobile device 580
  session 83
theHarvester 203, 725–726
thin clients 235–236, 508
“third countries” 220
third-party destruction and disposal services 387
third-party libraries 265
third-party risks 155–160
threat actors
  attack vectors 122–123
  attributes of 122
  types of 120–121
threat feeds 176
threat hunting 175–180
threat intelligence
  automated indicator sharing (AIS) 125
  Information Sharing and Analysis Centers (ISACs) 123–125
  MITRE ATT&CK framework 128–129
  research sources 127–128
  Structured Threat Information eXpression (STIX) 125–127
  Trusted Automated eXchange of Indicator Information (TAXII) 125–127
  vulnerability databases 125
threat maps 127
threat modeling 264
thumbprint algorithm 692

ticket-granting tickets (TGTs) 668
tickets, Kerberos 668
time 844–845
  delay 74
  offset 844
  synchronization 440
  time bombs 39
  time of check (TOC) attacks 79
  time of use (TOU) attacks 79
  time-based logins 639
  time-based one-time password (TOTP) 295
  timestamps 82–83, 844
Time Machine 158
time-to-live (TTL) 795
TKIP (Temporal Key Integrity Protocol) 552
TLS (Transport Layer Security) 82–83, 88, 108, 351, 410, 436, 441, 556, 577, 656, 698
TLSI (Transport Layer Security Inspection) 215–216
TMSAD (Trust Model for Security Automation Data) 887
TOC (time of check) attacks 79
token key 297
token-based authentication 297
tokenization 218, 461–462, 945
tokens 461, 625, 627–628
Top 10 Web Application Security Risks 277
Top Secret information 905, 941–942
TOS (trusted operating system) 160, 905
ToS (type of service) bits 536
Toshiba 476
TOTP (time-based one-time password) 295
TOU (time of use) attacks 79
TPM (Trusted Platform Module) 294, 459–460, 477–478, 524, 655
traceroute command 707–709
tracert command 707–709
traffic
  east-west 492
  lateral 492
training, user 22–24, 899, 901–902
Transaction Signature (TSIG) 108
transference of risk 918
transit gateways 246–247
transitive trust 577–578
Transmission Control Protocol (TCP) 503
transparent proxy 516
Transport Layer Security Inspection (TLSI) 215–216
Transport Layer Security (TLS) 82–83, 88, 108, 351, 410, 436, 441, 556, 577, 656, 698
transport mode, IPsec 438, 503
traversal, directory 75–76, 149, 274–275, 276
Triple DES 412
TRNG (true random number generators) 477
Trojans 35, 104, 108, 113
true random number generators (TRNGs) 477
trust
  models 698
  root of 476–477
  transitive 577–578
  Trusted Computer System Evaluation Criteria (TCSEC) 674
  zero 494
Trust Model for Security Automation Data (TMSAD) 887
Trusted Automated eXchange of Indicator Information (TAXII) 125–127
trusted computing base (TCB) 676
Trusted Computing Group (TCG) 476
trusted execution environment (TEE) 476
trusted operating system (TOS) 160, 905

Trusted Platform Module (TPM) 294, 459–460, 477–478, 524, 655
trusted zones 825
trustworthy computing 39–40
Try-SQL Editor 71
TSIG (Transaction Signature) 108
TTLS (Tunneled Transport Layer Security) 556
TTPs (tactics, techniques, and procedures) 128, 176, 767, 809
tunnel mode, IPsec 438, 503
Tunneled Transport Layer Security (TTLS) 556
tunneling 495–496, 505–508, 556
two-factor authentication (2FA) 298
Twofish 412
two-person integrity control 378
Type I errors 626
Type II errors 626
type of service (ToS) bits 536
typo squatting 20, 44

U

UAC (User Account Control) 67
UAs (user agents) 800
UAVs (unmanned aerial vehicles) 353–354
ubuntu keyword 239
UDP (User Datagram Protocol) 503
UEFI (Unified Extensible Firmware Interface) 459, 851
UEM (unified endpoint management) 587–588, 825
Umbrella 509
unauthenticated modes 404
unauthorized hackers 121
Unclassified information 941–942
underscore (_) 740
unicast addresses 537
unified endpoint management (UEM) 587–588, 825
Unified Extensible Firmware Interface (UEFI) 459, 851
unified threat management (UTM) 495, 524
uniform resource locators (URLs)
  redirection attacks 110
  URL hijacking 44
uninterruptible power source (UPS) 320–321
union operator 73
unique serial numbers (ICCIDs) 360
Universal Serial Bus. *See* USB (Universal Serial Bus)
UNIX 144
unknown environment 198–199
unmanned aerial vehicles (UAVs) 353–354
UPDATE statement 70
updates, exam 02.0004–02.0026
UPN (User Principal Name) 696
UPS (uninterruptible power source) 320–321
URLs (uniform resource locators)
  filtering 828–829
  redirection attacks 110
  URL hijacking 44
US Computer Emergency Readiness Team (US-CERT) 576
US Office of Personnel Management (OPM) attack 300–301
USB (Universal Serial Bus)
  condoms 379
  data blockers 379–380
  malicious flash drives 47–48
  malicious USB cables 48
  USB OTG (USB On-The-Go) 583
  USB sticks 123
US-DMCA (Digital Millennium Copyright Act) 220
use case analysis 882
user access recertification 645

User Account Control (UAC) 67
user accounts. *See* accounts
user agents (UAs) 800
user behavior analysis 188
user certificates 696
User Datagram Protocol (UDP) 503
user education 899, 901–902
USER environment variable 740
User Interaction (UI) metric 184
User Principal Name (UPN) 696
user security awareness education 22–24
user-controlled input 464
user-generated passwords 638
users command 631–632
UTC (Coordinated Universal Time) 845
UTM (unified threat management) 495, 524

V

vacations, mandatory 898–899, 900
validation
  continuous 278
  identity 633–635
  input 267–268, 464
validity dates, certificate 692
variables, environmental 740
/var/log directory 791
vaults 385, 655
VBA (Visual Basic for Applications) 113
VDEs (virtual desktop environments) 139, 232
VDIs (virtual desktop infrastructures) 139, 232
vectors, attack 122–123
vehicle systems 347–348
vein authentication 302
vendor management 155, 156, 331, 902–903
ver command 161
Veracode Web Application Security 204
Verisign 112, 577
version control 258, 279
vertical privilege escalation 67
vestibules, access control 372–373
video
  forensic video analysis 842–843
  secure 440
  steganography 416
virtualization 606, 610. *See also* VPNs (virtual private networks)
  APIs (application programming interfaces)
    definition of 240–241
    infrastructure as code 241–243
    micro-segmentation 240–241
  cloud computing
    cloud models 231–232
    cloud service providers (CSPs) 233
    community cloud 233
    fog and edge computing 234–235
    hybrid cloud 233
    managed detection and response (MDR) 234
    managed service providers (MSPs) 233–234
    off-premises versus on-premises services 234
    private cloud 232–233
    public cloud 232
    thin clients 235–236
    VPCs (virtual private clouds) 607, 608, 611
  containers 236–240
  definition of 247
  firewalls 534–535
  IP addresses 488
  memory 850
  microservices 236–240
  resource policies 246
  serverless architecture 243–244
  services integration 246
  transit gateways 246–247
  VDEs (virtual desktop environments) 139, 232

VDIs (virtual desktop infrastructures) 139, 232, 589
VLANs (virtual local-area networks) 490–491, 831
VMs (virtual machines) 324–325
attacks 248–249, 601
sprawl avoidance 247–248
VNC (virtual network computing) servers 632
VPCs (virtual private clouds) 607, 608, 611
viruses
antivirus software 451
fileless 37
vishing 12–13
visitor logs 383
Visual Basic for Applications (VBA) 113
VLANs (virtual local-area networks) 490–491, 831
VMs (virtual machines) 324–325
attacks 248–249, 601
sprawl avoidance 247–248
VNC (virtual network computing) servers 632
voice, secure 440
voice recognition 302
VoIP (Voice over Internet Protocol) 350, 799–800
volatility, order of 848
VPCs (virtual private clouds) 607, 608, 611
VPNs (virtual private networks) 99
always-on functionality 495
clientless versus client-based 497, 507
definition of 494
description of 494–496
example of 494–495
HTML5 508
IKEv1 Phase 1 negotiation 498–501
IKEv1 Phase 2 negotiation 501–503
IKEv2 504–505
IPsec 497, 501–502
Layer 2 Tunneling Protocol (L2TP) 508
remote-access 496–497
SCADA systems 341–342
site-to-site configuration 495, 496–497
split tunneling 495–496
SSL-based 505–508
VPN concentrators 495
vulnerabilities
backdoors 149, 271, 275
cloud-based versus on-premises 137–143
code injection 149, 273–274, 276
cross-site request forgery (XSRF) 149, 272, 275
cross-site scripting (XSS) 54, 68–70, 110, 149, 272, 275, 601
dark web 143
directory traversal 149, 274–275, 276
error handling 79–82
compile-time errors 81–82
input handling 79–82
runtime errors 81–82
impact of cybersecurity breaches and attacks 165–166
legacy platforms 165
memory/buffer 77–78, 149, 271–272, 275
peer to peer (P2P) networks 143
remote code execution (RCE) 78, 146, 149, 275
server defense 144
email servers 145
file servers 144
FTP servers 147–148
network controllers 144
web servers 146–147
social media 143
summary of 149–150, 275–276
third-party risks 155–160
vulnerability databases 125
weak configurations 150–155

weak patch management 160–164
zero-day 149, 275, 276, 522
vulnerability scans 180–181, 559
Common Vulnerability Scoring System (CVSS) 182–186
false negative 181
false positives 181
intrusive versus nonintrusive 182
noncredentialed 182
output 785–786
VUPEN Web Application Security Scanner 204

W

w command 631
WADL (Web Application Description Language) documents 87
WAF (web application firewall) 198, 531
walkthrough exercises 766
WannaCry 34, 37
WAP (Wireless Application Protocol) 558, 585
WAPs (wireless access points) 98, 101, 513, 559
war driving 205
war flying 205
war-dialing 13
warm sites 221
waterfall development methodology 257–258
watering hole attacks 20, 85
weak configurations 150–155
weak defaults 346
weak patch management 160–164
wearables 345
Web Application Description Language (WADL) documents 87
web application firewall (WAF) 198, 531
Web Application Proxy 516
web application scanners 182
Web form–grabbing keyloggers 43
web of trust 698
web pages, metadata from 808–809
web protocol port numbers 441
web servers
logs 794
vulnerabilities 146–147
Web Services Description Language (WSDL) documents 87
web synchronization 440
WebApp360 204
webification 507
Websense 533
WebSploit 72, 238, 953
weighted random early detection (WRED) 536
WEP (Wired Equivalent Privacy) 102
WER (Windows Error Reporting) 853
Western Digital 476
whaling 9, 16–17
white box testing 468–469
white hat hackers 121
white teams 206
white-box testing 80
WhiteHat Sentinel 204
whitelisting 578, 583
who command 631–632
whoami command 632
whois 108, 203
Wi-Fi
vulnerabilities and exposures 571
Wi-Fi ad hoc 584
Wi-Fi Analyzers 559, 561
Wi-Fi direct 584
Wi-Fi disassociation attack 101
WPA2 (Wi-Fi Protected Access 2) 551
WPA3 (Wi-Fi Protected Access 3) 551
WPS (Wi-Fi Protected Setup) 558–559
Wigle 205
wildcard certificates 694–695

Windows Defender Firewall 457
Windows Error Reporting (WER) 853
Windows Event Viewer 791–792, 846
Windows Performance Monitor 540–542
Windows Performance tool 539
Windows PowerShell 630
WinGate 514
WinHex 746
Wired Equivalent Privacy (WEP) 102
wireless access points (WAPs) 98, 101, 513, 559
Wireless Application Protocol (WAP) 558, 585
wireless LAN (WLAN) controllers 558
wireless networks 547, 557–558
  attacks 98, 122
    bluejacking 100
    bluesnarfing 99–100
    disassociation and deauthentication 101
    evil twin 98–99
    initialization vector (IV) 103
    jamming 102, 561–562
    mobile device security countermeasures 580
    near-field communication (NFC) 102–103
    radio frequency identification (RFID) 49, 102
    rogue access points 98–99
  authentication protocols 556–557
  cryptographic protocols 551
    Advanced Encryption Standard (AES) 552
    Counter-mode/CBC-MAC protocol (CCMP) 552
    Simultaneous Authentication of Equals (SAE) 551, 552
    summary of 552
    Wi-Fi Protected Access 2 (WPA2) 551
    Wi-Fi Protected Access 3 (WPA3) 551–552
  installation considerations
    AP isolation 562
    captive portals 559
    controller and access point security 562–563
    firewalls 562
    heat maps 559
    IEEE 802.1X standard 562
    Multi-User Multiple Input (MU-MIMO) 560–561
    orthogonal frequency-division multiple access (OFDMA) 561
    site surveys 559, 561–562
    Wi-Fi Analyzer tools 559
    Wi-Fi Protected Setup (WPS) 558–559
    wireless access point (WAP) placement 559
Wireless Transport Layer Security (WTLS) 558
Wireshark 539, 559, 743
WLAN (wireless LAN) controllers 558
workstations, hardening 159–160
WORM (write once read many) device 789
worms 36–37
WPA2 (Wi-Fi Protected Access 2) 551
WPA3 (Wi-Fi Protected Access 3) 551
WPS (Wi-Fi Protected Setup) 558–559
wrap 77
write once read many (WORM) devices 789
WSDL (Web Services Description Language) documents 87
WTLS (Wireless Transport Layer Security) 558
wuapp.exe 161

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

X

X.509 standard 694
X.690 encoding formats 697
XaaS (anything as a service) 139, 232
XCCDF (Extensible Configuration Checklist Description Format) 885
XDR (Extended Detection and Response) 189
Xiao 415
XML (Extensible Markup Language)
  XML injection 74–75, 273–274
  XSD (XML Schema Definition) 86
  XXE (XML External Entity) 74
XSRF (cross-site request forgery) 85–86, 149, 272, 275
XSS (cross-site scripting) 54, 68–70, 110, 149, 272, 275, 464, 601
Xways 850–851
X-Ways Software Technology AG 746
XXE (XML External Entity) 75

Y

YOLO (You Only Look Once) 376
YubiKey 297

Z

Zed Attack Proxy 204
zero trust 494
zero-day vulnerabilities 149, 275, 276, 522
Zigbee 360–361
Zimbra 145
zombies 111–112
zones
  high availability across 603, 609
  zone signing keys (ZSKs) 427
  zone transfers 109
ZSKs (zone signing keys) 427
Zune 850–851

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

This page intentionally left blank

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Exclusive Offer – 40% OFF

Pearson IT Certification Video Training

livelessons

pearsonitcertification.com/video

Use coupon code **PITCVIDEO40** during checkout.



Video Instruction from Technology Experts



Advance Your Skills

Get started with fundamentals, become an expert, or get certified.



Train Anywhere

Train anywhere, at your own pace, on any device.



Learn

Learn from trusted author trainers published by Pearson IT Certification.

Try Our Popular Video Training for FREE!

pearsonitcertification.com/video

Explore hundreds of **FREE** video lessons from our growing library of Complete Video Courses, LiveLessons, networking talks, and workshops.

PEARSON
IT CERTIFICATION

pearsonitcertification.com/video

ALWAYS LEARNING

PEARSON



Photo by Olena Yakobchuk/Shutterstock

Register Your Product at pearsonITcertification.com/register

Access additional benefits and **save 35%** on your next purchase

- Automatically receive a coupon for 35% off your next purchase, valid for 30 days. Look for your code in your Pearson IT Certification cart or the Manage Codes section of your account page.
- Download available product updates.
- Access bonus material if available.*
- Check the box to hear from us and receive exclusive offers on new editions and related products.

**Registration benefits vary by product. Benefits will be listed on your account page under Registered Products.*

Learning Solutions for Self-Paced Study, Enterprise, and the Classroom

Pearson IT Certification delivers training materials that address the learning, preparation, and practice needs of a new generation of certification candidates, including the official publishing programs of Adobe Press, Cisco Press, and Microsoft Press. At pearsonITcertification.com, you can:

- Shop our books, eBooks, practice tests, software, and video courses
- Sign up to receive special offers
- Access thousands of free chapters and video lessons

Visit **pearsonITcertification.com/community** to connect with Pearson IT Certification

Pearson

Addison-Wesley • Adobe Press • Cisco Press • Microsoft Press • Pearson IT Certification • Que • Sams • Peachpit Press

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

To receive your 10% off Exam Voucher, register your product at:

www.pearsonitcertification.com/register

and follow the instructions.

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Appendix C

Study Planner

Practice Test	Readir	Task

Element	Task	Goal Date	First Date Completed	Second Date Completed (Optional)	Notes
Introduction	Read Introduction				
1. Comparing and Contrasting Different Types of Social Engineering Techniques	Read Foundation Topics				
1. Comparing and Contrasting Different Types of Social Engineering Techniques	Review Key Topics				
1. Comparing and Contrasting Different Types of Social Engineering Techniques	Define Key Terms				
1. Comparing and Contrasting Different Types of Social Engineering Techniques	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 1 in practice test software				
2. Analyzing Potential Indicators to Determine the Type of Attack	Read Foundation Topics				
2. Analyzing Potential Indicators to Determine the Type of Attack	Review Key Topics				
2. Analyzing Potential Indicators to Determine the Type of Attack	Define Key Terms				
2. Analyzing Potential Indicators to Determine the Type of Attack	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 2 in practice test software				
3. Analyzing Potential Indicators Associated with Application Attacks	Read Foundation Topics				
3. Analyzing Potential Indicators Associated with Application Attacks	Review Key Topics				
3. Analyzing Potential Indicators Associated with Application Attacks	Define Key Terms				
3. Analyzing Potential Indicators Associated with Application Attacks	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 3 in practice test software				
4. Analyzing Potential Indicators Associated with Network Attacks	Read Foundation Topics				

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

4. Analyzing Potential Indicators Associated with Network Attacks	Review Key Topics				
4. Analyzing Potential Indicators Associated with Network Attacks	Define Key Terms				
4. Analyzing Potential Indicators Associated with Network Attacks	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 4 in practice test software				
5. Understanding Different Threat Actors, Vectors, and Intelligence Sources	Read Foundation Topics				
5. Understanding Different Threat Actors, Vectors, and Intelligence Sources	Review Key Topics				
5. Understanding Different Threat Actors, Vectors, and Intelligence Sources	Define Key Terms				
5. Understanding Different Threat Actors, Vectors, and Intelligence Sources	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 5 in practice test software				
6. Understanding the Security Concerns Associated with Various Types of Vulnerabilities	Read Foundation Topics				
6. Understanding the Security Concerns Associated with Various Types of Vulnerabilities	Review Key Topics				
6. Understanding the Security Concerns Associated with Various Types of Vulnerabilities	Define Key Terms				
6. Understanding the Security Concerns Associated with Various Types of Vulnerabilities	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 6 in practice test software				
7. Summarizing the Techniques Used in Security Assessments	Read Foundation Topics				
7. Summarizing the Techniques Used in Security Assessments	Review Key Topics				
7. Summarizing the Techniques Used in Security Assessments	Define Key Terms				
7. Summarizing the Techniques Used in Security Assessments	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 7 in practice test software				
8. Understanding the Techniques Used in Penetration Testing	Read Foundation Topics				
8. Understanding the Techniques Used in Penetration Testing	Review Key Topics				
8. Understanding the Techniques Used in Penetration Testing	Define Key Terms				
8. Understanding the Techniques Used in Penetration Testing	Complete Review Questions section				

Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 8 in practice test software				
9. Understanding the Importance of Security Concepts in an Enterprise Environment	Read Foundation Topics				
9. Understanding the Importance of Security Concepts in an Enterprise Environment	Review Key Topics				
9. Understanding the Importance of Security Concepts in an Enterprise Environment	Define Key Terms				
9. Understanding the Importance of Security Concepts in an Enterprise Environment	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 9 in practice test software				
10. Summarizing Virtualization and Cloud Computing Concepts	Read Foundation Topics				
10. Summarizing Virtualization and Cloud Computing Concepts	Review Key Topics				
10. Summarizing Virtualization and Cloud Computing Concepts	Define Key Terms				
10. Summarizing Virtualization and Cloud Computing Concepts	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 10 in practice test software				
11. Summarizing Secure Application Development, Deployment, and Automation Concepts	Read Foundation Topics				
11. Summarizing Secure Application Development, Deployment, and Automation Concepts	Review Key Topics				
11. Summarizing Secure Application Development, Deployment, and Automation Concepts	Define Key Terms				
11. Summarizing Secure Application Development, Deployment, and Automation Concepts	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 11 in practice test software				
12. Summarizing Authentication and Authorization Design Concepts	Read Foundation Topics				
12. Summarizing Authentication and Authorization Design Concepts	Review Key Topics				
12.Summarizing Authentication and Authorization Design Concepts	Define Key Terms				
12. Summarizing Authentication and Authorization Design Concepts	Complete Review Questions section				

Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 12 in practice test software				
13. Implementing Cybersecurity Resilience	Read Foundation Topics				
13. Implementing Cybersecurity Resilience	Review Key Topics				
13. Implementing Cybersecurity Resilience	Define Key Terms				
13. Implementing Cybersecurity Resilience	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 13 in practice test software				
14. Understanding the Security Implications of Embedded and Specialized Systems	Read Foundation Topics				
14. Understanding the Security Implications of Embedded and Specialized Systems	Review Key Topics				
14. Understanding the Security Implications of Embedded and Specialized Systems	Define Key Terms				
14. Understanding the Security Implications of Embedded and Specialized Systems	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 14 in practice test software				
15. Understanding the Importance of Physical Security Controls	Read Foundation Topics				
15. Understanding the Importance of Physical Security Controls	Review Key Topics				
15. Understanding the Importance of Physical Security Controls	Define Key Terms				
15. Understanding the Importance of Physical Security Controls	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 15 in practice test software				
16. Summarizing the Basics of Cryptographic Concepts	Read Foundation Topics				
16. Summarizing the Basics of Cryptographic Concepts	Review Key Topics				
16. Summarizing the Basics of Cryptographic Concepts	Define Key Terms				
16. Summarizing the Basics of Cryptographic Concepts	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 16 in practice test software				
17. Implementing Secure Protocols	Read Foundation Topics				
17. Implementing Secure Protocols	Review Key Topics				
17. Implementing Secure Protocols	Define Key Terms				

17. Implementing Secure Protocols	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 17 in practice test software				
18. Implementing Host or Application Security Solutions	Read Foundation Topics				
18. Implementing Host or Application Security Solutions	Review Key Topics				
18. Implementing Host or Application Security Solutions	Define Key Terms				
18. Implementing Host or Application Security Solutions	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 18 in practice test software				
19. Implementing Secure Network Designs	Read Foundation Topics				
19. Implementing Secure Network Designs	Review Key Topics				
19. Implementing Secure Network Designs	Define Key Terms				
19. Implementing Secure Network Designs	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 19 in practice test software				
20. Installing and Configuring Wireless Security Settings	Read Foundation Topics				
20. Installing and Configuring Wireless Security Settings	Review Key Topics				
20. Installing and Configuring Wireless Security Settings	Define Key Terms				
20. Installing and Configuring Wireless Security Settings	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 20 in practice test software				
21. Implementing Secure Mobile Solutions	Read Foundation Topics				
21. Implementing Secure Mobile Solutions	Review Key Topics				
21. Implementing Secure Mobile Solutions	Define Key Terms				
21. Implementing Secure Mobile Solutions	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 21 in practice test software				
22. Applying Cybersecurity Solutions to the Cloud	Read Foundation Topics				
22. Applying Cybersecurity Solutions to the Cloud	Review Key Topics				

22. Applying Cybersecurity Solutions to the Cloud	Define Key Terms				
22. Applying Cybersecurity Solutions to the Cloud	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 22 in practice test software				
23. Implementing Identity and Account Management Controls	Read Foundation Topics				
23. Implementing Identity and Account Management Controls	Review Key Topics				
23. Implementing Identity and Account Management Controls	Define Key Terms				
23. Implementing Identity and Account Management Controls	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 23 in practice test software				
24. Implementing Authentication and Authorization Solutions	Read Foundation Topics				
24. Implementing Authentication and Authorization Solutions	Review Key Topics				
24. Implementing Authentication and Authorization Solutions	Define Key Terms				
24. Implementing Authentication and Authorization Solutions	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 24 in practice test software				
25. Implementing Public Key Infrastructure	Read Foundation Topics				
25. Implementing Public Key Infrastructure	Review Key Topics				
25. Implementing Public Key Infrastructure	Define Key Terms				
25. Implementing Public Key Infrastructure	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 25 in practice test software				
26. Using the Appropriate Tool to Assess Organizational Security	Read Foundation Topics				
26. Using the Appropriate Tool to Assess Organizational Security	Review Key Topics				
26. Using the Appropriate Tool to Assess Organizational Security	Define Key Terms				
26. Using the Appropriate Tool to Assess Organizational Security	Complete Review Questions section				

Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 26 in practice test software				
27. Summarizing the Importance of Policies, Processes, and Procedures for Incident Response	Read Foundation Topics				
27. Summarizing the Importance of Policies, Processes, and Procedures for Incident Response	Review Key Topics				
27. Summarizing the Importance of Policies, Processes, and Procedures for Incident Response	Define Key Terms				
27. Summarizing the Importance of Policies, Processes, and Procedures for Incident Response	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 27 in practice test software				
28. Using Appropriate Data Sources to Support an Investigation	Read Foundation Topics				
28. Using Appropriate Data Sources to Support an Investigation	Review Key Topics				
28. Using Appropriate Data Sources to Support an Investigation	Define Key Terms				
28. Using Appropriate Data Sources to Support an Investigation	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 28 in practice test software				
29. Applying Mitigation Techniques or Controls to Secure an Environment	Read Foundation Topics				
29. Applying Mitigation Techniques or Controls to Secure an Environment	Review Key Topics				
29. Applying Mitigation Techniques or Controls to Secure an Environment	Define Key Terms				
29. Applying Mitigation Techniques or Controls to Secure an Environment	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 29 in practice test software				
30. Understanding the Key Aspects of Digital Forensics	Read Foundation Topics				
30. Understanding the Key Aspects of Digital Forensics	Review Key Topics				
30. Understanding the Key Aspects of Digital Forensics	Define Key Terms				
30. Understanding the Key Aspects of Digital Forensics	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 30 in practice test software				

31. Comparing and Contrasting the Various Types of Controls	Read Foundation Topics				
31. Comparing and Contrasting the Various Types of Controls	Review Key Topics				
31. Comparing and Contrasting the Various Types of Controls	Define Key Terms				
31. Comparing and Contrasting the Various Types of Controls	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 31 in practice test software				
32. Understanding the Importance of Applicable Regulations, Standards, or Frameworks That Impact Organizational Security Posture	Read Foundation Topics				
32. Understanding the Importance of Applicable Regulations, Standards, or Frameworks That Impact Organizational Security Posture	Review Key Topics				
32. Understanding the Importance of Applicable Regulations, Standards, or Frameworks That Impact Organizational Security Posture	Define Key Terms				
32. Understanding the Importance of Applicable Regulations, Standards, or Frameworks That Impact Organizational Security Posture	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 32 in practice test software				
33. Understanding the Importance of Policies to Organizational Security	Read Foundation Topics				
33. Understanding the Importance of Policies to Organizational Security	Review Key Topics				
33. Understanding the Importance of Policies to Organizational Security	Define Key Terms				
33. Understanding the Importance of Policies to Organizational Security	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 33 in practice test software				
34. Summarizing Risk Management Processes and Concepts	Read Foundation Topics				
34. Summarizing Risk Management Processes and Concepts	Review Key Topics				
34. Summarizing Risk Management Processes and Concepts	Define Key Terms				
34. Summarizing Risk Management Processes and Concepts	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 34 in practice test software				
35. Understanding Privacy and Sensitive Data Concepts in Relation to Security	Read Foundation Topics				

35. Understanding Privacy and Sensitive Data Concepts in Relation to Security	Review Key Topics				
35. Understanding Privacy and Sensitive Data Concepts in Relation to Security	Define Key Terms				
35. Understanding Privacy and Sensitive Data Concepts in Relation to Security	Complete Review Questions section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 35 in practice test software				
36. Final Preparation					
36. Final Preparation	Take practice test in study mode for all book questions in practice test software				
36. Final Preparation	Review all Key Topics in all chapters				
36. Final Preparation	Take practice test in practice exam mode using Exam Bank #1 questions for all chapters				
36. Final Preparation	Take practice test in practice exam mode using Exam Bank #2 questions for all chapters				

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

CompTIA Security+ SY0-601 Cert Guide **Companion Website**

Access interactive study tools on this book's companion website, including practice test software, memory table exercises, a Key Term flash card application, study planner, and more!

To access the companion website, simply follow these steps:

1. Go to **www.pearsonITcertification.com/register**.
2. Enter the print book ISBN: **9780136770312**.
3. Answer the security question to validate your purchase.
4. Go to your account page.
5. Click on the **Registered Products** tab.
6. Under the book listing, click on the Access Bonus Content link.

If you have any issues accessing the companion website, you can contact our support team by going to **http://pearsonitp.echelp.org.**

Humble Bundle Pearson Networking and Security Certification Bundle – © Pearson. Do Not Distribute.

Pearson

Where are the companion content files?



Register this digital version of
CompTIA Security+ SY0-601 Cert Guide
to access important downloads.

Register this eBook to unlock the companion files. Follow these steps:

1. Go to **pearsonITcertification.com/account** and log in or create a new account.
2. Enter the ISBN: **9780136770312** (NOTE: Please enter the print book ISBN provided to register the eBook you purchased.)
3. Answer the challenge question as proof of purchase.
4. Click on the "Access Bonus Content" link in the Registered Products section of your account page, to be taken to the page where your downloadable content is available.

This eBook version of the print title does not contain the practice test software that accompanies the print book.

You May Also Like—Premium Edition eBook and Practice Test. To learn about the Premium Edition eBook and Practice Test series, visit **pearsonITcertification.com/practicetest**

The Professional and Personal Technology Brands of Pearson

Addison Wesley Cisco Press informIT PEARSON IT Certification QUE SAMS

Humble Bundle Pearson Network and Security Certification Bundle. © Pearson. Do Not Distribute