


CompTIA

SECURITY+

Get Certified Get Ahead

SY0-401 Study Guide

- Real-world examples of security principles in action
- Over 400 realistic practice test questions with in-depth explanations
- 100 percent coverage of all CompTIA Security+ SY0-401 exam objectives
- Save 10 percent on your exam voucher
Access to coupon inside



Darril Gibson

CompTIA A+, Network+, Security+,
CASP, (ISC)2 SSCP, CISSP

CompTIA Security+:
Get Certified Get Ahead
SY0-401 Study Guide

Darril Gibson

Dedication

To my wife, who even after 22 years of marriage continues to remind me how wonderful life can be if you're in a loving relationship. Thanks for sharing your life with me.

Acknowledgments

Books of this size and depth can't be done by a single person, and I'm grateful for the many people who helped me put this book together. First, thanks to my wife. She has provided me immeasurable support throughout this project. The technical editor, Steve Johnson, provided some good feedback throughout the project. If you have the paperback copy of the book in your hand, you're enjoying some excellent composite editing work done by Susan Veach.

I'm extremely grateful for all the effort Karen Annett put into this project. She's an awesome copy editor and proofer and the book is tremendously better due to all the work she's put into it.

While I certainly appreciate all the feedback everyone gave me, I want to stress that any technical errors that may have snuck into this book are entirely my fault and no reflection on anyone who helped. I always strive to identify and remove every error, but they still seem to sneak in.

About the Author

Darril Gibson is the CEO of YCDA, LLC (short for You Can Do Anything). He has contributed to more than 35 books as the sole author, a coauthor, or a technical editor. Darril regularly writes, consults, and teaches on a wide variety of technical and security topics and holds several certifications, including CompTIA A+, Network+, Security+, and CASP; (ISC)2 SSCP and CISSP; Microsoft MCSE and MCITP, and ITIL Foundations.

In response to repeated requests, Darril created the <http://gcgapremium.com/> site where he provides study materials for several certification exams, including the Security+ exam. Darril regularly posts blog articles at <http://blogs.getcertifiedgetahead.com/>, and uses this site to help people stay abreast of changes in certification exams. You can contact him through either of these sites.

Darril lives in Virginia Beach with his wife and two dogs. Whenever possible, they escape to a small cabin in the country on over twenty acres of land that continue to provide them with peace, tranquility, and balance.

Table of Contents

[Dedication](#)

[Acknowledgments](#)

[About the Author](#)

[Introduction](#)

[*Who This Book Is For*](#)

[*About This Book*](#)

[*How to Use This Book*](#)

[*Remember This*](#)

[*Vendor Neutral*](#)

[*Web Resources*](#)

[*Assumptions*](#)

[*Set a Goal*](#)

[About the Exam](#)

[*Number of Questions and Duration*](#)

[*Passing Score*](#)

[*Exam Prerequisites*](#)

[*Exam Format*](#)

[*Beta Questions*](#)

[*Question Types*](#)

[Multiple Choice](#)

[Performance-Based Questions](#)

[*Exam Test Provider*](#)

[*Voucher Code for 10 Percent Off*](#)

[*Exam Domains*](#)

[*Objective to Chapter Map*](#)

[*Recertification Requirements*](#)

[CompTIA Security+ Assessment Exam](#)

[Assessment Exam Answers](#)

[Chapter 1 Mastering Security Basics](#)

[Understanding Core Security Goals](#)

[Confidentiality](#)

[Encryption](#)

[Access Controls](#)

[Steganography](#)

[Integrity](#)

[Hashing](#)

[Digital Signatures, Certificates, and Non-Repudiation](#)

[Availability](#)

[Redundancy and Fault Tolerance](#)

[Patching](#)

[Safety](#)

[Layered Security/Defense in Depth](#)

[Introducing Basic Risk Concepts](#)

[Exploring Authentication Concepts](#)

[Comparing Identification, Authentication, and Authorization](#)

[Verifying Identities with Identity Proofing](#)

[Identity Proofing for Verification](#)

[Self-Service Password Reset Systems](#)

[Comparing Authentication Factors](#)

[Something You Know](#)

[Something You Have](#)

[Something You Are](#)

[Somewhere You Are](#)

[Something You Do](#)

[Dual-Factor and Multifactor Authentication](#)

[Summarizing Identification Methods](#)

[Comparing Authentication Services](#)

[Kerberos](#)

[LDAP and Secure LDAP](#)

[Single Sign-On](#)

[SSO and Transitive Trusts](#)

[SSO and a Federation](#)

[SSO and SAML](#)

[SAML and Authorization](#)

[Authenticating RAS Clients](#)

[PAP](#)

[CHAP](#)

[MS-CHAP and MS-CHAPv2](#)

[RADIUS](#)

[Diameter](#)

[XTACACS](#)

[TACACS+](#)

[AAA Protocols](#)

[Chapter 1 Exam Topic Review](#)

[Chapter 1 Practice Questions](#)

[Chapter 1 Practice Question Answers](#)

[Chapter 2 Exploring Control Types and Methods](#)

[Understanding Control Types](#)

[Control Implementation Methods](#)

[Technical Controls](#)

[Management Controls](#)

[Operational Controls](#)

[Control Goals](#)

[Preventive Controls](#)

[Detective Controls](#)

[Comparing Detection and Prevention Controls](#)

[Corrective Controls](#)

[Deterrent Controls](#)

[Compensating Controls](#)

[Combining Control Types and Goals](#)

[Comparing Physical Security Controls](#)

[Comparing Door Access Systems](#)

[Securing Door Access with Cipher Locks](#)

[Securing Door Access with Proximity Cards](#)

[Securing Door Access with Biometrics](#)

[Identifying Users with ID Badges](#)

[Tailgating](#)

[Preventing Tailgating with Mantraps](#)

[Increasing Physical Security with Guards](#)

[Controlling Access with Access Lists and Logs](#)

[Monitoring Areas with Video Surveillance](#)

[Combining Fencing and Motion Detection](#)

[Combining Proper Lighting and Motion Detection](#)

[Combining Alarms and Motion Detection](#)

[Securing Access with Barricades](#)

[Using Signs](#)

[Using Hardware Locks](#)

[Securing Mobile Computers with Cable Locks](#)

[Securing Servers with Locking Cabinets](#)

[Securing Small Devices with a Safe](#)

[Implementing Logical Access Controls](#)

[Least Privilege](#)

[Need to Know](#)

[Group Policy](#)

[Using a Password Policy](#)

[Domain Password Policy](#)

[Application Passwords](#)

[Managing Accounts](#)

[Disabling and Deleting Accounts](#)

[Recovering Accounts](#)

[Prohibiting Generic Accounts](#)

[Restricting Access Based on Time-of-Day](#)

[Expiring Accounts](#)

[Reviewing Account Access](#)

[Credential Management](#)

[Comparing Access Control Models](#)

[Role-Based Access Control](#)

[Using Roles Based on Jobs and Functions](#)

[Documenting Roles with a Matrix](#)

[Establishing Access with Group-Based Privileges](#)

[Rule-Based Access Control](#)

[Discretionary Access Control](#)

[SIDs and DACLS](#)

[The Owner Establishes Access](#)

[Beware of Trojans](#)

[Mandatory Access Control](#)

[Labels and Lattice](#)

[Establishing Access](#)

[Chapter 2 Exam Topic Review](#)

[Chapter 2 Practice Questions](#)

[Chapter 2 Practice Question Answers](#)

[Chapter 3 Understanding Basic Network Security](#)

[Reviewing Basic Networking Concepts](#)

[Protocols](#)

[Common TCP/IP Protocols](#)

[IPv4](#)

[IPv6](#)

[Understanding DNS](#)

[Understanding and Identifying Ports](#)

[Combining the IP Address and the Port](#)

[IP Address Used to Locate Hosts](#)

[Server Ports](#)

[Client Ports](#)

[Putting It All Together](#)

[The Importance of Ports in Security](#)

[Memorize These Ports](#)

[Understanding Basic Network Devices](#)

[Hub](#)

[Switch](#)

[Security Benefit of a Switch](#)

[Physical Security of a Switch](#)

[Loop Protection](#)

[VLAN](#)

[Port Security](#)

[802.1x Port Security](#)

[Router](#)

[Routers and ACLs](#)

[Implicit Deny](#)

[Firewall](#)

[Host-Based Firewalls](#)

[Network-Based Firewalls](#)

[Firewall Rules](#)

[Web Application Firewall](#)

[Advanced Firewalls](#)

[Firewall Logs and Log Analysis](#)

[Network Separation](#)

[Protecting the Network Perimeter](#)

[DMZ](#)

[Understanding NAT and PAT](#)

[Proxies](#)

[Caching Content for Performance](#)

[Using URL Filters to Restrict Access](#)

[Unified Threat Management](#)

[Web Security Gateway](#)

[UTM Security Appliances](#)

[Web Security Gateway Versus UTM Security Appliance](#)

[Identifying OSI Relevance](#)

[Understanding the Layers](#)

[Layer 1: Physical](#)

[Layer 2: Data Link](#)

[Layer 3: Network](#)

[Layer 4: Transport](#)

[Layer 5: Session](#)

[Layer 6: Presentation](#)

[Layer 7: Application](#)

[Firewall Rules Solution](#)

[Chapter 3 Exam Topic Review](#)

[Chapter 3 Practice Questions](#)

[Chapter 3 Practice Question Answers](#)

[Chapter 4 Securing Your Network](#)

[Understanding IDSs and IPSs](#)

[***Packet Sniffing***](#)

[HIDS](#)

[NIDS](#)

[Detection Methods](#)

[Signature-Based Detection](#)

[Anomaly-Based Detection](#)

[Data Sources and Trends](#)

[Reporting](#)

[False Positives Versus False Negatives](#)

[IDS Responses](#)

[Honeypots](#)

[Honeynets](#)

[Counterattacks](#)

[IDS Versus IPS](#)

[Securing Wireless Networks](#)

[Reviewing Wireless Basics](#)

[WAPs and Wireless Routers](#)

[802.11](#)

[Antennas](#)

[Site Surveys and Antenna Placement](#)

[Security Protocols](#)

[WEP](#)

[WPA](#)

[WPA2](#)

[TKIP Versus CCMP](#)

[IEEE 802.1x](#)

[Personal Versus Enterprise Modes](#)

[EAP, PEAP, and LEAP](#)

[**WTLS and ECC**](#)

[Captive Portals](#)

[***Hot Spots and Isolation Mode***](#)

[Other Security Concerns](#)

[Change Default Administrator Password](#)

[Enable MAC Filtering](#)

[War Driving](#)

[War Biking](#)

[War Chalking](#)

[Change Default SSID](#)

[Disable SSID Broadcasting or Not](#)

[WEP/WPA Attacks](#)

[Rogue Access Points](#)

[Evil Twins](#)

[Jamming and Interference](#)

[Near Field Communication Attacks](#)

[Bluetooth Wireless](#)

[Discovery Mode](#)

[Bluetooth Attacks](#)

[Exploring Remote Access](#)

[Telephony](#)

[Dial-Up RAS](#)

[VPNs and VPN Concentrators](#)

[Connecting via Remote Access](#)

[IPsec as a Tunneling Protocol](#)

[TLS and SSL](#)

[PPTP](#)

[Site-to-Site VPNs](#)

[VPN Over Open Wireless](#)

[Network Access Control](#)

[Inspection and Control](#)

[**Chapter 4 Exam Topic Review**](#)

[**Chapter 4 Practice Questions**](#)

[**Chapter 4 Practice Question Answers**](#)

[Chapter 5 Securing Hosts and Data](#)

[Implementing Host Security](#)

[OS and Application Hardening](#)

[Disabling Unnecessary Services](#)

[Eliminating Unneeded Applications](#)

[Disabling Unnecessary Accounts](#)

[Protecting Management Interfaces and Applications](#)

[Using Baselines](#)

[Security Baselines](#)

[Configuration Baselines](#)

[Host Software Baselines](#)

[Application Configuration Baselines](#)

[Performance Baselines](#)

[Baseline Reporting](#)

[Whitelisting Versus Blacklisting Applications](#)

[Trusted OS](#)

[Understanding Virtualization](#)

[Snapshots](#)

[Sandboxing and Security Control Testing](#)

[VMs as Files](#)

[Networking Connectivity](#)

[Risks Associated with Virtualization](#)

[Implementing Patch Management](#)

[Automated Versus Controlled Deployment](#)

[Scheduling Patch Management](#)

[Testing Patches](#)

[Deploying and Verifying Patches](#)

[Mitigating Risk in Static Environments](#)

[Understanding Stuxnet](#)

[Protecting Static Systems](#)

[Securing Mobile Devices](#)

[Device Security](#)

[BYOD Concerns](#)

[Mobile Device Management](#)

[Application Security](#)

[Protecting Data](#)

[Comparing Data Categories](#)

[Protecting Confidentiality with Encryption](#)

[Software-Based Encryption](#)

[Hardware-Based Encryption](#)

[Data Leakage](#)

[Data Loss Prevention](#)

[Understanding SANs](#)

[Fibre Channel](#)

[iSCSI](#)

[FCoE](#)

[Handling Big Data](#)

[Understanding Cloud Computing](#)

[Software as a Service](#)

[Platform as a Service](#)

[Infrastructure as a Service](#)

[Public Versus Private Cloud](#)

[Cloud Computing Risks](#)

[Chapter 5 Exam Topic Review](#)

[Chapter 5 Practice Questions](#)

[Chapter 5 Practice Question Answers](#)

[Chapter 6 Understanding Malware and Social Engineering](#)

[Understanding Malware Types](#)

[Viruses](#)

[Armored Virus](#)

[Polymorphic Malware](#)

[Worms](#)

[Logic Bombs](#)

[Backdoors](#)

[Trojans](#)

[Botnets](#)

[Ransomware](#)

[The Police Virus](#)

[CryptoLocker](#)

[Rootkits](#)

[Spyware](#)

[Adware](#)

[Recognizing Common Attacks](#)

[Social Engineering](#)

[Impersonating](#)

[Shoulder Surfing](#)

[Tricking Users with Hoaxes](#)

[Tailgating and Mantraps](#)

[Dumpster Diving](#)

[Recognizing Other Attacks](#)

[Spam](#)

[Phishing](#)

[Spear Phishing](#)

[Whaling](#)

[Spim](#)

[Vishing](#)

[Privilege Escalation](#)

[Blocking Malware and Other Attacks](#)

[Protecting Systems with Anti-Malware Software](#)

[Antivirus Software](#)

[Pop-Up Blockers](#)

[Spam Filters as Anti-Spam Solutions](#)

[Anti-Spyware Software](#)

[Educating Users](#)

[New Viruses](#)

[Phishing Attacks](#)

[Zero-Day Exploits](#)

[Why Social Engineering Works](#)

[Authority](#)

[Intimidation](#)

[Consensus/Social Proof](#)

[Scarcity](#)

[Urgency](#)

[Familiarity/Liking](#)

[Trust](#)

[Chapter 6 Exam Topic Review](#)

[Chapter 6 Practice Questions](#)

[Chapter 6 Practice Question Answers](#)

[Chapter 7 Identifying Advanced Attacks](#)

[Comparing Common Attacks](#)

[Spoofing](#)

[DoS Versus DDoS](#)

[Smurf Attacks](#)

[SYN Flood Attacks](#)

[Flood Guards](#)

[Xmas Attacks](#)

[Man-in-the-Middle Attacks](#)

[Replay Attacks](#)

[Password Attacks](#)

[Brute Force Attacks](#)

[Dictionary Attacks](#)

[Password Hashes](#)

[Birthday Attacks](#)

[Rainbow Table Attacks](#)

[Hybrid Attacks](#)

[DNS Attacks](#)

[DNS Poisoning Attacks](#)

[Pharming Attacks](#)

[ARP Poisoning Attacks](#)

[ARP Man-in-the-Middle Attacks](#)

[ARP DoS Attack](#)

[Typo Squatting/URL Hijacking](#)

[Watering Hole Attacks](#)

[Zero-Day Attacks](#)

[Web Browser Concerns](#)

[Malicious Add-Ons](#)

[Cookies and Attachments](#)

[Session Hijacking Attacks](#)

[Flash Cookies and LSOs](#)

[Arbitrary Code Execution/Remote Code Execution](#)

[Header Manipulation Attacks](#)

[Understanding Secure Coding Concepts](#)

[Performing Input Validation](#)

[Client-Side and Server-Side Input Validation](#)

[Avoiding Race Conditions](#)

[Error and Exception Handling](#)

[Identifying Application Attacks](#)

[Web Servers](#)

[Buffer Overflows and Buffer Overflow Attacks](#)

[Integer Overflow](#)

[SQL Queries and SQL Injection Attacks](#)

[SQL Queries](#)

[SQL Injection Attacks](#)

[Protecting Against SQL Injection Attacks](#)

[XML Injection](#)

[NoSQL Versus SQL Databases](#)

[Cross-Site Scripting](#)

[Cross-Site Request Forgery \(XSRF\)](#)

[Directory Traversal/Command Injection](#)

[LDAP Injection](#)

[Transitive Access and Client-Side Attacks](#)

[Fuzzing](#)

[Chapter 7 Exam Topic Review](#)

[Chapter 7 Practice Questions](#)

[Chapter 7 Practice Question Answers](#)

[Chapter 8 Managing Risk](#)

[Identifying Risk](#)

[Threats and Threat Vectors](#)

[Types of Threats](#)

[Malicious Insider Threat](#)

[Threat Assessments](#)

[Vulnerabilities](#)

[Risk Management](#)

[Risk Assessment](#)

[Using Metrics to Identify Risk](#)

[Checking for Vulnerabilities](#)

[Anatomy of an Attack](#)

[Identifying IP Addresses of Targets](#)

[Identifying Open Ports with a Port Scanner](#)

[Fingerprint System](#)

[Banner Grabbing](#)

[Identifying Vulnerabilities](#)

[Attack](#)

[Putting It All Together](#)

[Vulnerability Assessment](#)

[Vulnerability Scanning](#)

[Other Assessment Techniques](#)

[Credentialed Versus Noncredentialed](#)

[Penetration Testing](#)

[White, Gray, and Black Box Testing](#)

[Obtaining Consent](#)

[Intrusive Versus Nonintrusive Testing](#)

[Passive Versus Active Tools](#)

[Continuous Monitoring](#)

[Identifying Security Tools](#)

[Sniffing with a Protocol Analyzer](#)

[Performing Routine Audits](#)

[User Reviews](#)

[Monitoring Events with Logs](#)

[Operating System Event Logs](#)

[Firewall and Router Access Logs](#)

[Other Logs](#)

[Reviewing Logs](#)

[Chapter 8 Exam Topic Review](#)

[Chapter 8 Practice Questions](#)

[Chapter 8 Practice Question Answers](#)

[Chapter 9 Preparing for Business Continuity](#)

[Adding Redundancy](#)

[Single Point of Failure](#)

[Disk Redundancies](#)

[RAID-0](#)

[RAID-1](#)

[RAID-5 and RAID-6](#)

[RAID-10](#)

[Software Versus Hardware RAID](#)

[Server Redundancy](#)

[Failover Clusters for High Availability](#)

[Load Balancers for High Availability](#)

[Power Redundancies](#)

[UPS](#)

[Generators](#)

[Protecting Data with Backups](#)

[Comparing Backup Types](#)

[Full Backups](#)

[Restoring a Full Backup](#)

[Differential Backups](#)

[Restoring a Full/Differential Backup Set](#)

[Incremental Backups](#)

[Restoring a Full/Incremental Backup Set](#)

[Choosing Full/Incremental or Full/Differential](#)

[Testing Backups](#)

[Protecting Backups](#)

[Backup Policies and Plans](#)

[Comparing Business Continuity Elements](#)

[Business Impact Analysis](#)

[Recovery Time Objective](#)

[Recovery Point Objective](#)

[Continuity of Operations](#)

[Hot Site](#)

[Cold Site](#)

[Warm Site](#)

[Site Variations](#)

[After the Disaster](#)

[Disaster Recovery](#)

[Planning for Communications](#)

[IT Contingency Planning](#)

[Succession Planning](#)

[BCP and DRP Testing](#)

[Testing Controls](#)

[Escape Plans, Escape Routes, and Drills](#)

[Implementing Environmental Controls](#)

[Heating, Ventilation, and Air Conditioning](#)

[Hot and Cold Aisles](#)

[HVAC and Fire](#)

[Fail-Safe Versus Fail-Open](#)

[Fire Suppression](#)

[Environmental Monitoring](#)

[Shielding](#)

[Shielding Cables](#)

[Protected Distribution of Cabling](#)

[Faraday Cage](#)

[Chapter 9 Exam Topic Review](#)

[Chapter 9 Practice Questions](#)

[Chapter 9 Practice Question Answers](#)

[Chapter 10 Understanding Cryptography](#)

[Introducing Cryptography Concepts](#)

[Providing Integrity with Hashing](#)

[MD5](#)

[SHA](#)

[HMAC](#)

[Hashing Files](#)

[Hashing Passwords](#)

[Hashing Messages](#)

[Using HMAC](#)

[Other Hash Algorithms](#)

[RIPEMD](#)

[LANMAN and NTLM](#)

[Providing Confidentiality with Encryption](#)

[Symmetric Encryption](#)

[Block Versus Stream Ciphers](#)

[AES](#)

[DES](#)

[3DES](#)

[RC4](#)

[Blowfish and Twofish](#)

[One-Time Pad](#)

[Asymmetric Encryption](#)

[The Rayburn Box](#)

[The Rayburn Box Used to Send Secrets](#)

[The Rayburn Box Used for Authentication](#)

[The Rayburn Box Demystified](#)

[Certificates](#)

[RSA](#)

[Static Versus Ephemeral Keys](#)

[Elliptic Curve Cryptography](#)

[Diffie-Hellman](#)

[Steganography](#)

[Quantum Cryptography](#)

[Using Cryptographic Protocols](#)

[Protecting Email](#)

[Signing Email with Digital Signatures](#)

[Encrypting Email](#)

[S/MIME](#)

[PGP/GPG](#)

[Transport Encryption](#)

[IPsec](#)

[SSL](#)

[TLS](#)

[Cipher Suites](#)

[Strong Versus Weak Ciphers](#)

[Encrypting HTTPS Traffic with SSL or TLS](#)

[Key Stretching](#)

[In-Band Versus Out-of-Band Key Exchange](#)

[Exploring PKI Components](#)

[Certificate Authority](#)

[Certificate Trust Paths and Trust Models](#)

[Self-Signed Certificates](#)

[Wildcard Certificates](#)

[Registration](#)

[Revoking Certificates](#)

[Validating Certificates](#)

[Outdated Certificates](#)

[Key Escrow](#)

[Recovery Agent](#)

[Chapter 10 Exam Topic Review](#)

[**Chapter 10 Practice Questions**](#)

[**Chapter 10 Practice Question Answers**](#)

[Chapter 11 Exploring Operational Security](#)

[Exploring Security Policies](#)

[Personnel Policies](#)

[Acceptable Use Policy and Privacy Policy](#)

[Mandatory Vacations](#)

[Separation of Duties](#)

[Job Rotation](#)

[Clean Desk Policy](#)

[Account Management Policies](#)

[Require Administrators to Use Two Accounts](#)

[Never Use Shared Accounts](#)

[Third-Party Issues](#)

[Interoperability Agreements](#)

[Change Management Policy](#)

[Data Policies](#)

[Information Classification](#)

[Data Labeling and Handling](#)

[Data Wiping and Disposing](#)

[Wiping Files](#)

[Storage and Retention Policies](#)

[Personally Identifiable Information](#)

[Protecting PII](#)

[Privacy Policy](#)

[Social Media Networks and Applications](#)

[Responding to Incidents](#)

[Incident Response Team](#)

[Incident Response Procedures](#)

[Implementing Basic Forensic Procedures](#)

[Order of Volatility](#)

[Capture System Image](#)

[Take Hashes](#)

[Network Traffic and Logs](#)

[Chain of Custody](#)

[Capture Video](#)

[Record Time Offset](#)

[Screenshots](#)

[Witnesses](#)

[Track Man-Hours and Expense](#)

[Big Data Analysis](#)

[Raising Security Awareness](#)

[Security Policy Training and Procedures](#)

[Role-Based Training](#)

[Training and Compliance Issues](#)

[Using Metrics to Validate Compliance](#)

[Chapter 11 Exam Topic Review](#)

[Chapter 11 Practice Questions](#)

[Chapter 11 Practice Question Answers](#)

[CompTIA Security+ Practice Exam](#)

[Security+ Practice Exam Answers](#)

[**Appendix A—Acronym List**](#)

Introduction

Congratulations on your purchase of *CompTIA Security+: Get Certified Get Ahead*. You are one step closer to becoming CompTIA Security+ certified. This certification has helped many individuals get ahead in their jobs and their careers, and it can help you get ahead, too.

It is a popular certification within the IT field. One IT hiring manager told me that if a résumé doesn't include the Security+ certification, or a higher-level security certification, he simply sets it aside. He won't even talk to applicants. That's not the same with all IT hiring managers, but it does help illustrate how important security is within the IT field.

Who This Book Is For

If you're studying for the CompTIA Security+ exam and want to pass it on your first attempt, this book is for you. It covers 100 percent of the objectives identified by CompTIA in enough depth so that you'll be able to easily answer the exam questions.

The first target audience for this book is students in CompTIA Security+ classes. My goal is to give students a book they can use to study the relevant and important details of CompTIA Security+ in adequate depth for the challenging topics, but without the minutiae in topics that are clear for most IT professionals. I regularly taught from the first edition of this book, and I'll continue to teach using this edition. I also heard from instructors around the United States and in other countries who used versions of the book to help students master the topics and pass the Security+ exam the first time they took it.

This book is also for those people who can study on their own. If you're one of the people who can read a book and learn the material without sitting in a class, this book has what you need to take and pass the exam the first time.

Additionally, you can keep this book on your shelf (or in your Kindle) to remind yourself of important, relevant concepts. These concepts are important for security professionals and IT professionals in the real world.

Based on many conversations with students and readers of the previous versions of this book, I know that many people use the Security+ certification as the first step in achieving other security certifications. For example, you may follow the Security+ with the ISC(2) SSCP or CISSP, or possibly the CompTIA CASP certification. If you plan to pursue any of these advanced security certifications, you'll find this book will help you lay a solid foundation of security knowledge. Learn this material, and you'll be a step ahead on the other exams.

About This Book

Over the past several years, I've taught literally hundreds of students, helping them to become CompTIA Security+ certified. During that time, I've learned what concepts are easy to grasp and what concepts need more explanation. I've developed handouts and analogies that help students grasp the elusive concepts.

Feedback from students was overwhelmingly positive—both in their comments to me and their successful pass rates after taking the certification exam. When the objectives changed in 2008, I rewrote my handouts as the first edition of this book. When the objectives changed again in 2011, I rewrote the book to reflect the new objectives. This book reflects the objective changes in 2014.

This book has allowed me to reach a much larger audience and share security and IT-related information. Even if you aren't in one of the classes I teach, this book can help you learn the relevant material to pass the exam the first time you take it.

How to Use This Book

When practicing for any certification exam, the following steps are a good recipe for success:

- **Review the objectives.** The objectives for the SY0-401 exam are listed in the “Objective to Chapter Map” in this Introduction.
- **Learn the material related to the objectives.** This book covers all of the objectives, and the introduction includes a map showing which chapter (or chapters) covers each objective.
- **Take practice questions.** A key step when preparing for any certification exam is to make sure you can answer the exam questions. Yes, you need the knowledge, but you also must be able to read a question and choose the correct answer. This simply takes practice. When using practice test questions, ensure they have explanations. Questions without explanations often give you the wrong answers.
- **Read and understand the explanations.** When preparing, you should make sure you know why the correct answers are correct and why the incorrect answers are incorrect. The explanations provide this information and are worded to help you get other questions correct.

This book has over 400 practice test questions you can use to test your knowledge and your ability to correctly answer them. Every question has a detailed explanation to help you understand why the correct answers are correct and why the incorrect answers are incorrect.

You can find the practice questions in the following areas:

- **Pre-assessment exam.** Use these questions at the beginning of the book to get a feel for what you know and what you need to study more.

- **End-of-chapter practice questions.** Each chapter has practice questions to help you test your comprehension of the material in the chapter.
- **End-of-book practice exam.** Use this as a practice exam to test your comprehension of the subject matter and readiness to take the actual exam.

It's OK if you do the practice questions in a different order. You may decide to tackle all the chapters in the book and then do the pre-assessment and post-assessment questions. That's fine. However, I strongly suggest you review all the questions in the book.

Remember This

Throughout the book, you'll see text boxes that highlight important information you should remember to successfully pass the exam. The surrounding content provides the additional information needed to fully understand these key points, and the text boxes summarize the important points.

These text boxes will look like this:

Remember this

I strongly encourage you to repeat the information in the text boxes to yourself as often as possible. The more you repeat the information, the more likely you are to remember it when you take the exam.

A tried-and-true method of repeating key information is to take notes when you're first studying the material and then rewrite the notes later. This will expose you to the material a minimum of three times.

Another method that students have told me has been successful for them is to use an MP3 player. Many MP3 players can record. Start your MP3 recorder and read the information in each text box for a chapter and the information in the Exam Topic Review section of each chapter. Save the MP3 file and regularly listen to it. This allows you to reaffirm the important information in your own voice.

You can play it while exercising, walking, or just about any time when it's not dangerous to listen to any MP3 file. You can even burn the MP3 files to a CD and play them back from a CD player.

If the MP3 method is successful for you, you can also record and listen to exam questions. Read the question, only the correct answer, and the first sentence or two of the explanation in each practice question.

If you don't have time to create your own MP3 recordings, check out the companion web site (<http://GetCertifiedGetAhead.com> and <http://gcgapremium.com>) for this book. You can purchase MP3 recordings there that you can download and use.

Vendor Neutral

CompTIA certifications are vendor neutral. In other words, certifications are not centered on any single vendor, such as Microsoft, Apple, or Linux. With that in mind, you don't need significantly deep knowledge of any of the operating systems, but don't be surprised if you see more questions about one OS over another simply because of market share.

In August 2014, Windows had about 91 percent market share of desktop and laptop computer operating systems. Apple MACs were next with about 6 percent, and Linux had about 1 percent. Looking at mobile operating systems such as tablets and smartphones, it's close to a tie with Android devices having 44.6 percent market share and Apple iOS systems having 44.19 percent.

Because over 90 percent of the systems you'll touch in a corporate environment are Microsoft based, don't be surprised to see some Microsoft-specific questions.

Web Resources

Check out <http://GetCertifiedGetAhead.com> for up-to-date details on the CompTIA Security+ exam. This site includes additional information related to the CompTIA Security+ exam and this book.

Although many people have spent a lot of time and energy trying to ensure that there are no errors in this book, occasionally they slip through. This site includes an errata page listing any errors we've discovered.

If you discover any errors, please let me know through the links on the web site. I'd also love to hear about your success when you pass the exam. I'm constantly getting good news from readers and students who are successfully earning their certifications.

In response to all the requests I've received for additional materials, such as online practice test questions, flash cards, and audio files, I created this site: <http://gcgapremium.com/>. It includes access to various study materials.

Last, I've found that many people find cryptography topics challenging, so I've posted some videos on YouTube (<http://www.youtube.com/>). As time allows, I'll post additional videos, and you can get a listing of all of them by searching YouTube with "Darril Gibson."

Assumptions

The CompTIA Security+ exam assumes you have at least two years of experience working with computers in a network. It also assumes you earned the CompTIA Network+ certification, or at least have the equivalent knowledge. While writing this book, I have largely assumed the same thing.

However, I'm well aware that two years of experience in a network could mean many different things. Your two years of experience may expose you to different technologies than someone else's two years of experience.

When it's critical that you understand an underlying network concept in order to master the relevant exam material, I have often included the concept within the background information.

Set a Goal

Look at a calendar right now and determine the date 45 days from today. This will be your target date to take this exam. Set this as your goal to complete studying the materials and to take the exam.

This target allows you to master about one and a half chapters per week. It may be that some of the chapters take you less time and some of the chapters take you more time. No problem. If you want to modify your target date later, do so. However, a recipe for success in almost any endeavor includes setting a goal.

When I teach CompTIA Security+ at a local university, I often help the students register for the exam on the first night. They pick a date close to the end of the course and register. I've found that when we do this, about 90 percent of the students take and pass the exam within one week after completing the course. On the other hand, when I didn't help the students register on the first night, more than half of them did not complete the exam in the same time frame. Setting a goal helps.

About the Exam

CompTIA first released the Security+ exam in 2002, and it has quickly grown in popularity. They revised the exam objectives in 2008, 2011, and again in 2014. The 2014 exam is numbered as SY0-401 (or JK0-022 for the academic version of the exam). SY0-201 retired on December 31, 2011, and SY0-301 is scheduled to retire on December 31, 2014.

The SY0-401 exam is the same as the JK0-022 exam. CompTIA uses the JK0-022 code for CompTIA Academy Partners. If you attend a Security+ course at a CompTIA Academy partner, they might give you a JK0-022 voucher. Everyone else uses the SY0-401 code.

A summary of the details of the exam includes:

- **Number of questions:** Maximum of 90 questions
- **Time to complete questions:** 90 minutes (does not include time to complete pretest and posttest surveys)
- **Passing score:** 750
- **Grading criteria:** Scale of 100 to 900 (about 83 percent)
- **Question types:** Multiple choice and performance-based
- **Exam format:** Traditional—can move back and forth to view previous questions
- **Exam prerequisites:** None required but Network+ is recommended
- **Exam test provider:** Pearson Vue

Number of Questions and Duration

You have 90 minutes to complete up to 90 questions. This gives you about one minute per question. Don't let this scare you; it's actually a good thing. With only about a minute to read and answer a question, you know the questions can't be very long. The exception is the performance-based questions, but you'll only see a few of those.

Passing Score

A score of 750 is required to pass. This is on a scale of 100 to 900. If the exam is paid for and you don't get a single question correct, you still get a score of 100. If you get every testable question correct, you get a score of 900.

If all questions are weighted equally, then you need to get 75 questions correct—a passing score of 750 divided by 900 equals .8333 or 83.33 percent. CompTIA doesn't say if all questions are scored equally or whether harder questions are weighted and worth more. However, most people believe that the performance-based questions are worth more than a typical multiple-choice question.

Also, a score of 83 percent is higher than many other certification exams, so you shouldn't underestimate the difficulty of this exam. However, many people regularly pass it and you can pass it, too. With this book, you will be well prepared.

Exam Prerequisites

All that is required for you to take the exam is money. Other than that, there are no enforced prerequisites. However, to successfully pass the exam, you're expected to have at least two years of experience working with computers in a networking environment. If you have more than that, the exam

materials will likely come easier to you. If you have less, the exam may be more difficult.

Exam Format

Questions are multiple-choice types where you choose one answer or multiple answers. When you need to choose multiple answers, the question may direct you to choose two, choose three, or choose all that apply.

You start at question 1 and go to the last question. During the process, you can mark any questions you want to review when you’re done. Additionally, you can view previous questions if desired. For example, if you get to question 10 and then remember something that helps you answer question 5, you can go back and redo question 5.

Beta Questions

Your exam may have some beta questions. They aren’t graded but instead are used to test the validity of the questions. If everyone gets a beta question correct, it’s probably too easy. If everyone gets it incorrect, there’s probably something wrong with the question. After enough people have tested a beta question, CompTIA personnel analyze it and decide if they want to add it to the test bank, or rewrite and test it as a new beta question.

The good news is that CompTIA doesn’t grade the beta questions. However, you don’t know which questions are ungraded beta questions and which questions are live questions, so you need to treat every question equally.

Question Types

Expect many of the questions on the exam to be straightforward. For example, what’s 5×5 ? Either you know the answer is 25 or you don’t. The exam questions test your knowledge of the material, not necessarily your ability to dissect the question so that you can figure out what the question is really trying to ask.

I’m not saying the knowledge is simplistic, only that the questions will be worded so that you can easily understand what they are asking.

As a comparative example, Microsoft certification questions can be quite complex. Microsoft questions often aren’t just testing your knowledge of the topic, but your ability to analyze the material and logically come to the right conclusion.

Here are two examples of questions—the first shows how Microsoft may word the question on a Microsoft certification exam, and the second shows how CompTIA may word it for the CompTIA Security+ exam.

- **Microsoft.** You are driving a bus from Chicago to Atlanta at 55 mph with 22 passengers. The bus is painted blue. At the same time, a train is traveling from Miami to Atlanta at 40 mph. The train has a yellow caboose. What color are the bus driver’s eyes?
- **CompTIA Security+.** What color are your eyes?

Notice the first question adds a lot of superfluous information. Two pieces are critical to answering the first question. It starts by saying, “You are driving a bus…” and then ends by asking, “What color are the bus driver’s eyes?” You’re required to put the two together and weed through the irrelevant information to come to the correct answer.

The second question is straightforward. “What color are your eyes?” There’s very little analysis required. Either you know it or you don’t. This is what you can expect from many of the CompTIA Security+ questions.

Some of the CompTIA exam questions may have a little more detail than just a single sentence, but overall, expect them to be one- to three-sentence questions. They are only giving you about one minute for each question, and it’s not intended to be a reading comprehension exam.

As a simple example, you may see a question like: “What port does SSH use?” In this case, you’d need to know that Secure Shell (SSH) uses Transmission Control Protocol (TCP) port 22.

However, CompTIA can reword the question to test your depth of comprehension. For example, you might see a question like this: “You need to configure a firewall to block Telnet traffic and allow traffic used by a more secure replacement of Telnet. What port needs to be opened on the firewall?” In this case, you’d need to know that Secure Shell (SSH) is a secure replacement for Telnet and SSH uses TCP port 22.

You may also see questions that use phrases such as “BEST choice,” “BEST description,” or “MOST secure.” In these examples, don’t be surprised if you see two answers that could answer the question, while only one is the best choice. For example, which one of the following numbers is between 1 and 10 and is the HIGHEST: 2, 8, 14, 23.

Clearly, 2 and 8 are between 1 and 10, but 14 and 23 are not. However, only 8 is both between 1 and 10 and the highest.

Here is a more realistic, security-related question that shows this:

Question: You need to send several large files containing proprietary data to a business partner. Which of the following is the BEST choice for this task?

- A. FTP
- B. SNMP
- C. SFTP

D. SSH

File Transfer Protocol (FTP) is a good choice to send large files. However, the question also says that the files include proprietary data, indicating they should be protected with encryption. Secure File Transfer Protocol (SFTP) is the best choice because it can send large files in an encrypted format. When you see key words like BEST or MOST, be careful not to jump on the first answer. There may be a more correct answer.

Multiple Choice

Most questions are multiple-choice types where you choose one answer or multiple answers. When you need to choose multiple answers, the question will include a phrase such as “Choose TWO” or “Choose THREE.”

Performance-Based Questions

You can expect as many as 10 non-multiple choice questions. CompTIA refers to these as performance-based questions and instead of picking from a multiple-choice answer, you’re often required to perform a task. CompTIA’s goal is to provide more accurate testing to verify people have a full understanding of a topic.

A question people often ask about these questions is if they get partial credit. People at CompTIA know, but I haven’t seen anywhere they’ve clarified this. It’s entirely possible that you get partial credit on some of these types of questions while others require you to answer them completely. It’s best to do the best you can with each question.

The following sections covers the different types of questions you can expect. You can also check out some of the blogs on performance-based questions that I’ve written here: <http://blogs.getcertifiedgetahead.com/security-blog-links/>.

Matching

In a matching performance-based question, you will see two lists and need to match them. As a simple example, one list might include several protocols and the other list might include several protocol numbers. You would need to match the correct protocol numbers with each protocol. If you know the protocols and ports listed in Table 3.1, this becomes trivial. Then again, if you don’t know which ports go with which protocols, this can be quite difficult.

Drag and Drop

In some questions, you might need to drag items from one location on the screen to another location to answer a question. You can think of these as multiple-choice questions with multiple answers that are correct. However, instead of selecting the check boxes to indicate a correct answer, you drag it to somewhere else on the screen.

Chapter 2, “Exploring Control Types and Methods,” has an example of a potential drag-and-drop question in the practice test questions section. It provides a scenario with several locations and devices, and lists several security controls used to secure the locations. You’re asked to drag and drop the appropriate controls to the appropriate location or device.

Another example is in Chapter 11, “Exploring Operational Security.” The question presents a scenario related to forensics, and then asks you to arrange an out-of-order list of data based on volatility.

Data Entry

Some performance-based questions might ask you to analyze a scenario and then enter appropriate data. For example, Chapter 4, “Securing Your Network,” discusses the configuration of wireless access points and wireless routers. A related question might ask you to configure an access point to work with WPA2 Enterprise mode. The Configuring a Wireless Router Lab mentioned in Chapter 4 and available online (<http://gcgapremium.com/labs/>) shows you the steps to do this.

Similarly, I wrote a series of blog articles on creating rules for routers and firewalls. The second post showed an example of a performance-based question and the last post provided the solution. You can read the posts here:

- **ACLs and Security+**. <http://blogs.getcertifiedgetahead.com/acls-and-security/>
- **Firewall Rules and Security+**. <http://blogs.getcertifiedgetahead.com/firewall-rules-and-security/>
- **Firewall Rules Solution**. <http://blogs.getcertifiedgetahead.com/firewall-rules-solution/>

Performance-Based Questions Strategy

You’ll see the performance-based questions first and they take much longer than typical multiple-choice questions. If the answer is clear to you, then by all means, take the time to answer it. However, if the question isn’t clear, mark the question and skip it. You can come back to it later. It’s entirely possible that the question is a poorly worded beta question that doesn’t even count. However, if you spend 45 minutes on it, you might run out of time before you finish the multiple-choice questions.

Performance-based questions have occasionally caused problems for the test systems. A common problem is that instead of displaying the question, the screen is mostly blank. If this happens, you can often just use the reset button for the question. This allows you to move past the problem and continue with the test. However, resetting the question erases any answer you’ve entered.

It’s common for people to be nervous when thinking about these performance-based test questions. However, the majority of people who take the test say that these questions really aren’t that difficult. As long as you understand the concepts from the exam objectives, you won’t have any problem. I do recommend you check out the posts on performance-based questions that I’ve posted here: <http://blogs.getcertifiedgetahead.com/security-blog-links/>.

Exam Test Provider

You can take the exam at a Pearson Vue testing site. Some testing sites provide testing and nothing else. However, most testing sites are part of another company, such as a training company, college, or university. You can take an exam at the training company’s testing site even if you haven’t taken a course with them.

The Pearson Vue web site includes search tools you can use to find a testing site close to you. Check them out at <http://www.pearsonvue.com>.

Voucher Code for 10 Percent Off

As of this writing, the CompTIA Security+ exam is \$293 in the United States if you purchase it at full price. However, you can get a 10 percent discount using a discount code. This code changes periodically, so you’ll need to go to this page to access the current code: <http://gcgapremium.com/discounted-comptia-vouchers/>.

When you purchase a voucher, you’ll get a voucher number that you can use to register at a testing site. A word of caution: Some criminals sell bogus vouchers on Internet sites such as eBay. You won’t know you’ve been ripped off until you try to use it and by that time, the criminal will probably have disappeared. In contrast, if you use the discount code, you buy the voucher directly from CompTIA.

Exam Domains

The exam objectives are divided into the following domains, or general topic areas. Additionally, CompTIA publishes the percentage of questions you can anticipate in any of the domains:

- **1.0 Network Security.** 20 percent of examination content
- **2.0 Compliance and Operational Security.** 18 percent of examination content
- **3.0 Threats and Vulnerabilities.** 20 percent of examination content
- **4.0 Application, Data, and Host Security.** 15 percent of examination content
- **5.0 Access Control and Identity Management.** 15 percent of examination content
- **6.0 Cryptography.** 12 percent of examination content

CompTIA publishes a listing of the objectives on its web site. As of this writing, this listing is accurate, but CompTIA includes the following disclaimers:

- “The lists of examples provided in bulleted format below each objective are not exhaustive lists. Other examples of technologies, processes or tasks pertaining to each objective may also be included on the exam although not listed or covered in this objectives document.”
- “The CompTIA Security+ Certification Exam Objectives are subject to change without notice.”

You can verify that the objectives haven’t changed by checking CompTIA’s web site (<http://www.comptia.org>). Additionally, you can check this book’s companion site at <http://GetCertifiedGetAhead.com> for up-to-date information on the exam, and read blogs about various topics, including the Security+ exam here: <http://blogs.GetCertifiedGetAhead.com>. Also, online practice test questions, flash cards, and other study materials are available here: <http://gcgapremium.com/>.

Objective to Chapter Map

This following listing shows the SY0-401 objectives published by CompTIA. In parentheses following the objective, you can also see the chapter or chapters where the objective is covered within this book.

1.0 Network Security

1.1 Implement security configuration parameters on network devices and other technologies. (Chapters 3, 4, 6, 8, 9)

- Firewalls (Chapter 3)
- Routers (Chapter 3)
- Switches (Chapter 3)
- Load Balancers (Chapter 9)

- Proxies (Chapter 3)
- Web security gateways (Chapter 3)
- VPN concentrators (Chapter 4)
- NIDS and NIPS (Chapter 4)
 - Behavior based (Chapter 4)
 - Signature based (Chapter 4)
 - Anomaly based (Chapter 4)
 - Heuristic (Chapter 4)
- Protocol analyzers (Chapter 8)
- Spam filter (Chapter 6)
- UTM security appliances (Chapter 3)
 - URL filter (Chapter 3)
 - Content inspection (Chapter 3)
 - Malware inspection (Chapter 3)
- Web application firewall vs. network firewall (Chapter 3)
- Application aware devices (Chapter 3)
 - Firewalls (Chapter 3)
 - IPS (Chapter 4)
 - IDS (Chapter 4)
 - Proxies (Chapter 3)

1.2 Given a scenario, use secure network administration principles. (Chapters 3, 4, 7)

- Rule-based management (Chapter 3)
- Firewall rules (Chapter 3)
- VLAN management (Chapter 3)
- Secure router configuration (Chapter 3)
- Access control lists (Chapter 3)
- Port Security (Chapter 3)
- 802.1x (Chapters 3, 4)
- Flood guards (Chapter 7)
- Loop protection (Chapter 3)
- Implicit deny (Chapter 3)
- Network separation (Chapter 3)
- Log analysis (Chapter 3)
- Unified Threat Management (Chapter 3)

1.3 Explain network design elements and components. (Chapters 1, 3, 4, 5)

- DMZ (Chapter 3)
- Subnetting (Chapter 3)
- VLAN (Chapter 3)
- NAT (Chapter 3)
- Remote Access (Chapter 4)
- Telephony (Chapter 4)
- NAC (Chapter 4)
- Virtualization (Chapter 5)
- Cloud Computing (Chapter 5)
 - Platform as a Service (Chapter 5)
 - Software as a Service (Chapter 5)
 - Infrastructure as a Service (Chapter 5)
 - Private (Chapter 5)
 - Public (Chapter 5)
 - Hybrid (Chapter 5)
 - Community (Chapter 5)
- Layered security / Defense in depth (Chapter 1)

1.4 Given a scenario, implement common protocols and services. (Chapters 3, 4, 5, 10)

- Protocols (Chapters 3, 4, 5, 10)
 - IPsec (Chapters 3, 4)
 - SNMP (Chapter 3)
 - SSH (Chapter 3)
 - DNS (Chapter 3)
 - TLS (Chapters 3, 4, 10)
 - SSL (Chapters 3, 4, 10)
 - TCP/IP (Chapter 3)
 - FTPS (Chapter 3)
 - HTTPS (Chapters 3, 10)

- SCP (Chapter 3)
 - ICMP (Chapter 3)
 - IPv4 (Chapter 3)
 - IPv6 (Chapter 3)
 - iSCSI (Chapter 5)
 - Fibre Channel (Chapter 5)
 - FCoE (Chapter 5)
 - FTP (Chapter 3)
 - SFTP (Chapter 3)
 - TFTP (Chapter 3)
 - TELNET (Chapter 3)
 - HTTP (Chapter 3)
 - NetBIOS (Chapter 3)
- Ports (Chapter 3)
 - 21 (Chapter 3)
 - 22 (Chapter 3)
 - 25 (Chapter 3)
 - 53 (Chapter 3)
 - 80 (Chapter 3)
 - 110 (Chapter 3)
 - 139 (Chapter 3)
 - 143 (Chapter 3)
 - 443 (Chapter 3)
 - 3389 (Chapter 3)
- OSI relevance (Chapter 3)

1.5 Given a scenario, troubleshoot security issues related to wireless networking. (Chapter 4)

- WPA (Chapter 4)
- WPA2 (Chapter 4)
- WEP (Chapter 4)
- EAP (Chapter 4)
- PEAP (Chapter 4)
- LEAP (Chapter 4)
- MAC filter (Chapter 4)
- Disable SSID broadcast (Chapter 4)
- TKIP (Chapter 4)
- CCMP (Chapter 4)
- Antenna Placement (Chapter 4)
- Power level controls (Chapter 4)
- Captive portals (Chapter 4)
- Antenna types (Chapter 4)
- Site surveys (Chapter 4)
- VPN (over open wireless) (Chapter 4)

2.0 Compliance and Operational Security

2.1 Explain the importance of risk related concepts. (Chapters 2, 4, 5, 8, 9, 11)

- Control types (Chapter 2)
 - Technical (Chapter 2)
 - Management (Chapter 2)
 - Operational (Chapter 2)
- False positives (Chapter 4)
- False negatives (Chapter 4)
- Importance of policies in reducing risk (Chapters 2, 11)
 - Privacy policy (Chapter 11)
 - Acceptable use (Chapter 11)
 - Security policy (Chapter 11)
 - Mandatory vacations (Chapter 11)
 - Job rotation (Chapter 11)
 - Separation of duties (Chapter 11)
 - Least privilege (Chapters 2, 11)
- Risk calculation (Chapter 8)
 - Likelihood (Chapter 8)
 - ALE (Chapter 8)
 - Impact (Chapter 8)
 - SLE (Chapter 8)
 - ARO (Chapter 8)

- MTTR (Chapter 8)
 - MTTF (Chapter 8)
 - MTBF (Chapter 8)
- Quantitative vs. qualitative (Chapter 8)
- Vulnerabilities (Chapter 8)
- Threat vectors (Chapter 8)
- Probability / threat likelihood (Chapter 8)
- Risk-avoidance, transference, acceptance, mitigation, deterrence (Chapter 8)
- Risks associated with Cloud Computing and Virtualization (Chapter 5)
- Recovery time objective and recovery point objective (Chapter 9)

2.2 Summarize the security implications of integrating systems and data with third parties. (Chapter 11)

- On-boarding/off-boarding business partners (Chapter 11)
- Social media networks and/or applications (Chapter 11)
- Interoperability agreements (Chapter 11)
 - SLA (Chapter 11)
 - BPA (Chapter 11)
 - MOU (Chapter 11)
 - ISA (Chapter 11)
- Privacy considerations (Chapter 11)
- Risk awareness (Chapter 11)
- Unauthorized data sharing (Chapter 11)
- Data ownership (Chapter 11)
- Data backups (Chapter 11)
- Follow security policy and procedures (Chapter 11)
- Review agreement requirements to verify compliance and performance standards (Chapter 11)

2.3 Given a scenario, implement appropriate risk mitigation strategies. (Chapters 2, 5, 8, 9, 11)

- Change management (Chapter 11)
- Incident management (Chapter 11)
- User rights and permissions reviews (Chapters 8, 11)
- Perform routine audits (Chapters 2, 8)
- Enforce policies and procedures to prevent data loss or theft (Chapters 5, 9, 11)
- Enforce technology controls (Chapter 5)
 - Data Loss Prevention (DLP) (Chapter 5)

2.4 Given a scenario, implement basic forensic procedures. (Chapter 11)

- Order of volatility (Chapter 11)
- Capture system image (Chapter 11)
- Network traffic and logs (Chapter 11)
- Capture video (Chapter 11)
- Record time offset (Chapter 11)
- Take hashes (Chapter 11)
- Screenshots (Chapter 11)
- Witnesses (Chapter 11)
- Track man hours and expense (Chapter 11)
- Chain of custody (Chapter 11)
- Big Data analysis (Chapter 11)

2.5 Summarize common incident response procedures. (Chapter 11)

- Preparation (Chapter 11)
- Incident identification (Chapter 11)
- Escalation and notification (Chapter 11)
- Mitigation steps (Chapter 11)
- Lessons learned (Chapter 11)
- Reporting (Chapter 11)
- Recovery/reconstitution procedures (Chapter 11)
- First responder (Chapter 11)
- Incident isolation (Chapter 11)
 - Quarantine (Chapter 11)
 - Device removal (Chapter 11)
- Data breach (Chapter 11)
- Damage and loss control (Chapter 11)

2.6 Explain the importance of security related awareness and training. (Chapters 1, 2, 5, 6, 11)

- Security policy training and procedures (Chapter 11)
- Role-based training (Chapter 11)
- Personally identifiable information (Chapter 11)
- Information classification (Chapter 11)
 - High (Chapter 11)
 - Medium (Chapter 11)
 - Low (Chapter 11)
 - Confidential (Chapter 11)
 - Private (Chapter 11)
 - Public (Chapter 11)
- Data labeling, handling and disposal (Chapter 11)
- Compliance with laws, best practices and standards (Chapter 11)
- User habits (Chapters 1, 5, 6, 11)
 - Password behaviors (Chapter 1)
 - Data handling (Chapter 11)
 - Clean desk policies (Chapter 11)
 - Prevent tailgating (Chapters 2, 6)
 - Personally owned devices (Chapter 5)
- New threats and new security trends/alerts (Chapter 6)
 - New viruses (Chapter 6)
 - Phishing attacks (Chapter 6)
 - Zero-day exploits (Chapter 6)
- Use of social networking and P2P (Chapter 11)
- Follow up and gather training metrics to validate compliance and security posture (Chapter 11)

2.7 Compare and contrast physical security and environmental controls. (Chapters 2, 9)

- Environmental controls (Chapter 9)
 - HVAC (Chapter 9)
 - Fire suppression (Chapter 9)
 - EMI shielding (Chapter 9)
 - Hot and cold aisles (Chapter 9)
 - Environmental monitoring (Chapter 9)
 - Temperature and humidity controls (Chapter 9)
- Physical security(Chapters 2, 9)
 - Hardware locks (Chapter 2)
 - Mantraps (Chapter 2)
 - Video Surveillance (Chapter 2)
 - Fencing (Chapter 2)
 - Proximity readers (Chapter 2)
 - Access list (Chapter 2)
 - Proper lighting (Chapter 2)
 - Signs (Chapter 2)
 - Guards (Chapter 2)
 - Barricades (Chapter 2)
 - Biometrics (Chapter 2)
 - Protected distribution (cabling) (Chapter 9)
 - Alarms (Chapter 2)
 - Motion detection (Chapter 2)
- Control types (Chapter 2)
 - Deterrent (Chapter 2)
 - Preventive (Chapter 2)
 - Detective (Chapter 2)
 - Compensating (Chapter 2)
 - Technical (Chapter 2)
 - Administrative (Chapter 2)

2.8 Summarize risk management best practices. (Chapters 8, 9)

- Business continuity concepts (Chapter 9)
 - Business impact analysis (Chapter 9)
 - Identification of critical systems and components (Chapter 9)
 - Removing single points of failure (Chapter 9)
 - Business continuity planning and testing (Chapter 9)
 - Risk assessment (Chapter 8)
 - Continuity of operations (Chapter 9)

- Disaster recovery (Chapter 9)
 - IT contingency planning (Chapter 9)
 - Succession planning (Chapter 9)
 - High availability (Chapter 9)
 - Redundancy (Chapter 9)
 - Tabletop exercises (Chapter 9)
- Fault tolerance (Chapter 9)
 - Hardware (Chapter 9)
 - RAID (Chapter 9)
 - Clustering (Chapter 9)
 - Load balancing (Chapter 9)
 - Servers (Chapter 9)
- Disaster recovery concepts (Chapter 9)
 - Backup plans/policies (Chapter 9)
 - Backup execution/frequency (Chapter 9)
 - Cold site (Chapter 9)
 - Hot site (Chapter 9)
 - Warm site (Chapter 9)

2.9 Given a scenario, select the appropriate control to meet the goals of security. (Chapters 1, 2, 5, 9, 10)

- Confidentiality (Chapters 1, 10)
 - Encryption (Chapters 1, 10)
 - Access controls (Chapter 1)
 - Steganography (Chapters 1, 10)
- Integrity (Chapters 1, 10)
 - Hashing (Chapters 1, 10)
 - Digital signatures (Chapters 1, 10)
 - Certificates (Chapters 1, 10)
 - Non-repudiation (Chapters 1, 10)
- Availability (Chapters 1, 5, 9)
 - Redundancy (Chapters 1, 9)
 - Fault tolerance (Chapters 1, 9)
 - Patching (Chapters 1, 5)
- Safety (Chapters 1, 2, 9)
 - Fencing (Chapters 1, 2)
 - Lighting (Chapters 1, 2)
 - Locks (Chapters 1, 2)
 - CCTV (Chapters 1, 2)
 - Escape plans (Chapters 1, 9)
 - Drills (Chapters 1, 9)
 - Escape routes (Chapters 1, 9)
 - Testing controls (Chapters 1, 9)

3.0 Threats and Vulnerabilities

3.1 Explain types of malware. (Chapter 6)

- Adware (Chapter 6)
- Virus (Chapter 6)
- Spyware (Chapter 6)
- Trojan (Chapter 6)
- Rootkits (Chapter 6)
- Backdoors (Chapter 6)
- Logic bomb (Chapter 6)
- Botnets (Chapter 6)
- Ransomware (Chapter 6)
- Polymorphic malware (Chapter 6)
- Armored virus (Chapter 6)

3.2 Summarize various types of attacks. (Chapters 6, 7, 8)

- Man-in-the-middle (Chapter 7)
- DDoS (Chapter 7)
- DoS (Chapter 7)
- Replay (Chapter 7)
- Smurf attack (Chapter 7)

- Spoofing (Chapter 7)
- Spam (Chapter 6)
- Phishing (Chapter 6)
- Spim (Chapter 6)
- Vishing (Chapter 6)
- Spear phishing (Chapter 6)
- Xmas attack (Chapter 7)
- Pharming (Chapter 7)
- Privilege escalation (Chapter 6)
- Malicious insider threat (Chapter 8)
- DNS poisoning and ARP poisoning (Chapter 7)
- Transitive access (Chapter 7)
- Client-side attacks (Chapter 7)
- Password attacks (Chapter 7)
 - Brute force (Chapter 7)
 - Dictionary attacks (Chapter 7)
 - Hybrid (Chapter 7)
 - Birthday attacks (Chapter 7)
 - Rainbow tables (Chapter 7)
- Typo squatting/URL hijacking (Chapter 7)
- Watering hole attack (Chapter 7)

3.3 Summarize social engineering attacks and the associated effectiveness with each attack. (Chapter 6)

- Shoulder surfing (Chapter 6)
- Dumpster diving (Chapter 6)
- Tailgating (Chapter 6)
- Impersonation (Chapter 6)
- Hoaxes (Chapter 6)
- Whaling (Chapter 6)
- Vishing (Chapter 6)
- Principles (reasons for effectiveness) (Chapter 6)
 - Authority (Chapter 6)
 - Intimidation (Chapter 6)
 - Consensus/Social proof (Chapter 6)
 - Scarcity (Chapter 6)
 - Urgency (Chapter 6)
 - Familiarity/liking (Chapter 6)
 - Trust (Chapter 6)

3.4 Explain types of wireless attacks. (Chapters 4, 7)

- Rogue access points (Chapter 4)
- Jamming/Interference (Chapter 4)
- Evil twin (Chapter 4)
- War driving (Chapter 4)
- Bluejacking (Chapter 4)
- Bluesnarfing (Chapter 4)
- War chalking (Chapter 4)
- IV attack (Chapter 4)
- Packet sniffing (Chapter 4)
- Near field communication (Chapter 4)
- Replay attacks (Chapter 7)
- WEP/WPA attacks (Chapter 4)
- WPS attacks (Chapter 4)

3.5 Explain types of application attacks. (Chapters 4, 7)

- Cross-site scripting (Chapter 7)
- SQL injection (Chapter 7)
- LDAP injection (Chapter 7)
- XML injection (Chapter 7)
- Directory traversal/command injection (Chapter 7)
- Buffer overflow (Chapter 7)
- Integer overflow (Chapter 7)
- Zero-day (Chapters 4, 7)

- Cookies and attachments (Chapter 7)
- LSO (Locally Shared Objects) (Chapter 7)
- Flash Cookies (Chapter 7)
- Malicious add-ons (Chapter 7)
- Session hijacking (Chapter 7)
- Header manipulation (Chapter 7)
- Arbitrary code execution / remote code execution (Chapter 7)

3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques. (Chapters 1, 2, 3, 4, 5, 8)

- Monitoring system logs (Chapter 8)
 - Event logs (Chapter 8)
 - Audit logs (Chapter 8)
 - Security logs (Chapter 8)
 - Access logs (Chapter 8)
- Hardening (Chapters 1, 5)
 - Disabling unnecessary services (Chapter 5)
 - Protecting management interfaces and applications (Chapter 5)
 - Password protection (Chapters 1, 5)
 - Disabling unnecessary accounts (Chapter 5)
- Network security (Chapters 3, 4)
 - MAC limiting and filtering (Chapters 3, 4)
 - 802.1x (Chapters 3, 4)
 - Disabling unused interfaces and unused application service ports (Chapter 3)
 - Rogue machine detection (Chapter 4)
- Security posture (Chapters 4, 5)
 - Initial baseline configuration (Chapter 5)
 - Continuous security monitoring (Chapter 4)
 - Remediation (Chapter 4)
- Reporting (Chapter 4)
 - Alarms (Chapter 4)
 - Alerts (Chapter 4)
 - Trends (Chapter 4)
- Detection controls vs. prevention controls (Chapters 2, 4)
 - IDS vs. IPS (Chapter 4)
 - Camera vs. guard (Chapter 2)

3.7 Given a scenario, use appropriate tools and techniques to discover security threats and vulnerabilities. (Chapters 3, 4, 8)

- Interpret results of security assessment tools (Chapter 8)
- Tools (Chapters 4, 8)
 - Protocol analyzer (Chapter 8)
 - Vulnerability scanner (Chapter 8)
 - Honeypots (Chapter 4)
 - Honeynets (Chapter 4)
 - Port scanner (Chapters 3, 8)
 - Passive vs. active tools (Chapter 8)
 - Banner grabbing (Chapter 8)
- Risk calculations (Chapter 8)
 - Threat vs. likelihood (Chapter 8)
- Assessment types (Chapter 8)
 - Risk (Chapter 8)
 - Threat (Chapter 8)
 - Vulnerability (Chapter 8)
- Assessment technique (Chapter 8)
 - Baseline reporting (Chapter 8)
 - Code review (Chapter 8)
 - Determine attack surface (Chapter 8)
 - Review architecture (Chapter 8)
 - Review designs (Chapter 8)

3.8 Explain the proper use of penetration testing versus vulnerability scanning. (Chapter 8)

- Penetration testing (Chapter 8)
 - Verify a threat exists (Chapter 8)
 - Bypass security controls (Chapter 8)

- Actively test security controls (Chapter 8)
 - Exploiting vulnerabilities (Chapter 8)
- Vulnerability scanning (Chapter 8)
 - Passively testing security controls (Chapter 8)
 - Identify vulnerability (Chapter 8)
 - Identify lack of security controls (Chapter 8)
 - Identify common misconfigurations (Chapter 8)
 - Intrusive vs. non-intrusive (Chapter 8)
 - Credentialed vs. non-credentialed (Chapter 8)
 - False positive (Chapter 8)
- Black box (Chapter 8)
- White box (Chapter 8)
- Gray box (Chapter 8)

4.0 Application, Data and Host Security

4.1 Explain the importance of application security controls and techniques. (Chapters 5, 7)

- Fuzzing (Chapter 7)
- Secure coding concepts (Chapter 7)
 - Error and exception handling (Chapter 7)
 - Input validation (Chapter 7)
- Cross-site scripting prevention (Chapter 7)
- Cross-site Request Forgery (XSRF) prevention (Chapter 7)
- Application configuration baseline (proper settings) (Chapter 5)
- Application hardening (Chapter 5)
- Application patch management (Chapter 5)
- NoSQL databases vs. SQL databases (Chapter 7)
- Server-side vs. Client-side validation (Chapter 7)

4.2 Summarize mobile security concepts and technologies. (Chapter 5)

- Device security (Chapter 5)
 - Full device encryption (Chapter 5)
 - Remote wiping (Chapter 5)
 - Lockout (Chapter 5)
 - Screen-locks (Chapter 5)
 - GPS (Chapter 5)
 - Application control (Chapter 5)
 - Storage segmentation (Chapter 5)
 - Asset tracking (Chapter 5)
 - Inventory control (Chapter 5)
 - Mobile device management (Chapter 5)
 - Device access control (Chapter 5)
 - Removable storage(Chapter 5)
 - Disabling unused features (Chapter 5)
- Application security (Chapter 5)
 - Key management (Chapter 5)
 - Credential management (Chapter 5)
 - Authentication (Chapter 5)
 - Geo-tagging (Chapter 5)
 - Encryption (Chapter 5)
 - Application whitelisting (Chapter 5)
 - Transitive trust/authentication (Chapter 5)
- BYOD concerns (Chapter 5)
 - Data ownership (Chapter 5)
 - Support ownership (Chapter 5)
 - Patch management (Chapter 5)
 - Antivirus management (Chapter 5)
 - Forensics (Chapter 5)
 - Privacy (Chapter 5)
 - On-boarding/off-boarding (Chapter 5)
 - Adherence to corporate policies (Chapter 5)
 - User acceptance (Chapter 5)
 - Architecture/infrastructure considerations (Chapter 5)

- Legal concerns (Chapter 5)
- Acceptable use policy (Chapter 5)
- On-board camera/video (Chapter 5)

4.3 Given a scenario, select the appropriate solution to establish host security. (Chapters 2, 3, 4, 5, 6)

- Operating system security and settings (Chapter 5)
- OS hardening (Chapter 5)
- Anti-malware (Chapter 6)
 - Antivirus (Chapter 6)
 - Anti-spam (Chapter 6)
 - Anti-spyware (Chapter 6)
 - Pop-up blockers (Chapter 6)
- Patch management (Chapter 5)
- White listing vs. black listing applications (Chapter 5)
- Trusted OS (Chapter 5)
- Host-based firewalls (Chapter 3)
- Host-based intrusion detection (Chapter 4)
- Hardware security (Chapter 2)
 - Cable locks (Chapter 2)
 - Safe (Chapter 2)
 - Locking cabinets (Chapter 2)
- Host software baselining (Chapter 5)
- Virtualization (Chapter 5)
 - Snapshots (Chapter 5)
 - Patch compatibility (Chapter 5)
 - Host availability/elasticity (Chapter 5)
 - Security control testing (Chapter 5)
 - Sandboxing (Chapter 5)

4.4 Implement the appropriate controls to ensure data security. (Chapters 5, 11)

- Cloud storage (Chapter 5)
- SAN (Chapter 5)
- Handling Big Data (Chapter 5)
- Data encryption (Chapter 5)
 - Full disk (Chapter 5)
 - Database (Chapter 5)
 - Individual files (Chapter 5)
 - Removable media (Chapter 5)
 - Mobile devices (Chapter 5)
- Hardware based encryption devices (Chapter 5)
 - TPM (Chapter 5)
 - HSM (Chapter 5)
 - USB encryption (Chapter 5)
 - Hard drive (Chapter 5)
- Data in-transit, Data at-rest, Data in-use (Chapter 5)
- Permissions/ACL (Chapter 5)
- Data policies (Chapter 11)
 - Wiping (Chapter 11)
 - Disposing (Chapter 11)
 - Retention (Chapter 11)
 - Storage (Chapter 11)

4.5 Compare and contrast alternative methods to mitigate security risks in static environments. (Chapter 5)

- Environments (Chapter 5)
 - SCADA (Chapter 5)
 - Embedded (Printer, Smart TV, HVAC control) (Chapter 5)
 - Android (Chapter 5)
 - iOS (Chapter 5)
 - Mainframe (Chapter 5)
 - Game consoles (Chapter 5)
 - In-vehicle computing systems (Chapter 5)
- Methods (Chapter 5)
 - Network segmentation (Chapter 5)
 - Security layers (Chapter 5)

- Application firewalls (Chapter 5)
- Manual updates (Chapter 5)
- Firmware version control (Chapter 5)
- Wrappers (Chapter 5)
- Control redundancy and diversity (Chapter 5)

5.0 Access Control and Identity Management

5.1 Compare and contrast the function and purpose of authentication services. (Chapter 1)

- RADIUS (Chapter 1)
- TACACS+ (Chapter 1)
- Kerberos (Chapter 1)
- LDAP (Chapter 1)
- XTACACS (Chapter 1)
- SAML (Chapter 1)
- Secure LDAP (Chapter 1)

5.2 Given a scenario, select the appropriate authentication, authorization or access control. (Chapters 1, 2, 3, 5, 11)

- Identification vs. authentication vs. authorization (Chapter 1)
- Authorization (Chapters 2, 3)
 - Least privilege (Chapters 2, 11)
 - Separation of duties (Chapter 11)
 - ACLs (Chapter 3)
 - Mandatory access (Chapter 2)
 - Discretionary access (Chapter 2)
 - Rule-based access control (Chapter 2)
 - Role-based access control (Chapter 2)
 - Time of day restrictions (Chapter 2)
- Authentication (Chapter 1)
 - Tokens (Chapter 1)
 - Common access card (Chapter 1)
 - Smart card (Chapter 1)
 - Multifactor authentication (Chapter 1)
 - TOTP (Chapter 1)
 - HOTP (Chapter 1)
 - CHAP (Chapter 1)
 - PAP (Chapter 1)
 - Single sign-on (Chapter 1)
 - Access control (Chapters 1, 2)
 - Implicit deny (Chapter 3)
 - Trusted OS (Chapter 5)
- Authentication factors (Chapter 1)
 - Something you are (Chapter 1)
 - Something you have (Chapter 1)
 - Something you know (Chapter 1)
 - Somewhere you are (Chapter 1)
 - Something you do (Chapter 1)
- Identification (Chapter 1)
 - Biometrics (Chapter 1)
 - Personal identification verification card (Chapter 1)
 - Username (Chapter 1)
- Federation (Chapter 1)
- Transitive trust/authentication (Chapter 1)

5.3 Install and configure security controls when performing account management, based on best practices. (Chapters 1, 2, 8, 11)

- Mitigate issues associated with users with multiple account/roles and/or shared accounts (Chapter 11)
- Account policy enforcement (Chapters 1, 2)
 - Credential management (Chapter 2)
 - Group policy (Chapter 2)
 - Password complexity (Chapters 1, 2)
 - Expiration (Chapter 2)
 - Recovery (Chapter 1)
 - Disablement (Chapter 2)
 - Lockout (Chapter 1)
 - Password history (Chapter 1)
 - Password reuse (Chapter 1)

- Password length (Chapter 1)
 - Generic account prohibition (Chapter 2)
- Group based privileges (Chapter 2)
- User assigned privileges (Chapter 2)
- User access reviews (Chapter 8)
- Continuous monitoring (Chapter 8)

6.0 Cryptography

6.1 Given a scenario, utilize general cryptography concepts. (Chapter 10)

- Symmetric vs. asymmetric (Chapter 10)
- Session keys (Chapter 10)
- In-band vs. out-of-band key exchange (Chapter 10)
- Fundamental differences and encryption methods (Chapter 10)
 - Block vs. stream (Chapter 10)
- Transport encryption (Chapter 10)
- Non-repudiation (Chapter 10)
- Hashing (Chapter 10)
- Key escrow (Chapter 10)
- Steganography (Chapter 10)
- Digital signatures (Chapter 10)
- Use of proven technologies (Chapter 10)
- Elliptic curve and quantum cryptography (Chapter 10)
- Ephemeral key (Chapter 10)
- Perfect forward secrecy (Chapter 10)

6.2 Given a scenario, use appropriate cryptographic methods. (Chapters 1, 3, 4, 10)

- WEP vs. WPA/WPA2 and preshared key (Chapter 4)
- MD5 (Chapter 10)
- SHA (Chapter 10)
- RIPEMD (Chapter 10)
- AES (Chapter 10)
- DES (Chapter 10)
- 3DES (Chapter 10)
- HMAC (Chapter 10)
- RSA (Chapter 10)
- Diffie-Hellman (Chapter 10)
- RC4 (Chapter 10)
- One-time pads (Chapter 10)
- NTLM (Chapter 10)
- NTLMv2 (Chapter 10)
- Blowfish (Chapter 10)
- PGP/GPG (Chapter 10)
- TwoFish (Chapter 10)
- DHE (Chapter 10)
- ECDHE (Chapter 10)
- CHAP (Chapter 1)
- PAP (Chapter 1)
- Comparative strengths and performance of algorithms (Chapter 10)
- Use of algorithms/protocols with transport encryption (Chapters 3, 4, 10)
 - SSL (Chapters 3, 4, 10)
 - TLS (Chapters 3, 4, 10)
 - IPSec (Chapters 3, 4, 10)
 - SSH (Chapters 3, 10)
 - HTTPS (Chapter 3, 10)
- Cipher suites (Chapter 10)
 - Strong vs. weak ciphers (Chapter 10)
- Key stretching (Chapter 10)
 - PBKDF2 (Chapter 10)
 - Bcrypt (Chapter 10)

6.3 Given a scenario, use appropriate PKI, certificate management and associated components. (Chapter 10)

- Certificate authorities and digital certificates (Chapter 10)
 - CA (Chapter 10)

- CRLs (Chapter 10)
 - OCSP (Chapter 10)
 - CSR (Chapter 10)
- PKI (Chapter 10)
- Recovery agent (Chapter 10)
- Public key (Chapter 10)
- Private key (Chapter 10)
- Registration (Chapter 10)
- Key escrow (Chapter 10)
- Trust models (Chapter 10)

Recertification Requirements

The CompTIA Security+ certification was previously a lifetime certification. You passed the exam once and you were certified for life. However, for anyone taking the exam after January 1, 2011, the certification expires after three years unless it is renewed.

You can renew the certification by either taking the next version of the exam or by enrolling in CompTIA's new Continuing Education (CE) program. You will be required to pay an annual fee of \$49 and earn a minimum of 50 Continuing Education Units (CEUs). You can earn CEUs through a variety of activities. Some examples include presenting or teaching topics to others, attending training sessions, participating in industry events or seminars, or writing relevant articles, white papers, blogs, or books.

For full details, check out the CompTIA web site: <https://certification.comptia.org/continuing-education>.

In response to several requests, I am creating courses that people can use to earn CEUs. I currently have one course approved by CompTIA and I expect to have more available. They will be available through the <http://gcgapremium.com/> site.

It Pays to Get Certified

In a digital world, digital literacy is an essential survival skill.

Certification demonstrates that you have the knowledge and skill to solve technical or business problems in virtually any business environment. CompTIA certifications are highly- valued credentials that qualify you for jobs, increased compensation and promotion.








Certification Helps Your Career

- Security is one of the highest demand job categories** – growing in importance as the frequency and severity of security threats continues to be a major concern for organizations around the world.
- Jobs for security administrators are expected to increase by 18%** -the skill set required for these types of jobs maps to the CompTIA Security+ certification.
- Network Security Administrators** - can earn as much as \$106,000 per year.
- CompTIA Security+ is the first step** - in starting your career as a Network Security Administrator or Systems Security Administrator.
- More than ¼ million** – individuals worldwide are CompTIA Security+ certified.
- CompTIA Security+ is regularly used in organizations** - such as Hitachi Systems, Fuji Xerox, HP, Dell, and a variety of major U.S. government contractors.
- Approved by the U.S. Department of Defense (DoD)** - as one of the required certification options in the DoD 8570.01-M directive, for Information Assurance Technical Level II and Management Level I job roles.

Steps to Getting Certified and Staying Certified

Review Exam Objectives	Review the Certification objectives to make sure you know what is covered in the exam http://certification.comptia.org/examobjectives.aspx
Practice for the Exam	<ul style="list-style-type: none"> After you have studied for the certification, review and answer the sample questions to get an idea what type of questions might be on the exam. http://certification.comptia.org/samplequestions.aspx
Purchase an Exam Voucher	Purchase exam vouchers on the CompTIA Marketplace. www.comptiastore.com
Take the Test!	Go to the Pearson VUE website and schedule a time to take your exam. http://www.pearsonvue.com/comptia/
Stay Certified! Continuing Education	<ul style="list-style-type: none"> Effective January 1, 2011, new CompTIA Security+ certifications are valid for three years from the date of certification. There are a number of ways the certification can be renewed. For more information go to: http://certification.comptia.org/ce

How to obtain more information

- **Visit CompTIA online** - <http://certification.comptia.org/home.aspx> to learn more about getting CompTIA certified.
- **Contact CompTIA** - call 866-835-8020 and choose Option 2 or email questions@comptia.org.
- **Connect with us :**     

CompTIA Security+ Assessment Exam

Use this assessment exam to test your knowledge of the topics before you start reading the book, and again before you take the live exam. An answer key with explanations is available at the end of the assessment exam.

1. A security administrator is implementing a security program that addresses confidentiality and availability. Of the following choices, what else should the administrator include?
- A. Ensure critical systems provide uninterrupted service.
 - B. Protect data in transit from unauthorized disclosure.
 - C. Ensure systems are not susceptible to unauthorized changes.
 - D. Secure data to prevent unauthorized disclosure.
2. You need to transmit PII via email and you want to maintain its confidentiality. Of the following choices, what is the BEST solution?
- A. Use hashes.
 - B. Encrypt it before sending.
 - C. Protect it with a digital signature.
 - D. Use RAID.
3. Lisa manages network devices in your organization and maintains copies of the configuration files for all the managed routers and switches. On a weekly basis, she creates hashes for these files and compares them with hashes she created on the same files the previous week. Which security goal is she pursuing?
- A. Confidentiality
 - B. Integrity
 - C. Availability
 - D. Safety
4. An organization wants to provide protection against malware attacks. Administrators have installed antivirus software on all computers. Additionally, they implemented a firewall and an IDS on the network. Which of the following BEST identifies this principle?
- A. Implicit deny
 - B. Layered security
 - C. Least privilege
 - D. Flood guard
5. Homer called into the help desk and says he forgot his password. Which of the following choices is the BEST choice for what the help-desk professional should do?
- A. Verify the user’s account exists.
 - B. Look up the user’s password and tell the user what it is.
 - C. Disable the user’s account.
 - D. Reset the password and configure the password to expire after the first use.
6. Which type of authentication does a hardware token provide?
- A. Biometric
 - B. PIN
 - C. Strong password
 - D. One-time password
7. Which type of authentication is a retina scan?
- A. Multifactor
 - B. TOTP
 - C. Biometric
 - D. Dual-factor

8. Users are required to log on to their computers with a smart card and a PIN. Which of the following BEST describes this?
- A. Single-factor authentication
 - B. Multifactor authentication
 - C. Mutual authentication
 - D. TOTP
9. Your company recently began allowing workers to telecommute from home one or more days a week. However, your company doesn't currently have a remote access solution. They want to implement an AAA solution that supports different vendors. Which of the following is the BEST choice?
- A. TACACS+
 - B. RADIUS
 - C. Circumference
 - D. SAML
10. Your organization has implemented a system that stores user credentials in a central database. Users log on once with their credentials. They can then access other systems in the organization without logging on again. What does this describe?
- A. Same sign-on
 - B. SAML
 - C. Single sign-on
 - D. Biometrics
11. Your organization issues users a variety of different mobile devices. However, management wants to reduce potential data losses if the devices are lost or stolen. Which of the following is the BEST technical control to achieve this goal?
- A. Cable locks
 - B. Risk assessment
 - C. Disk encryption
 - D. Hardening the systems
12. Your primary job activities include monitoring security logs, analyzing trend reports, and installing CCTV systems. Which of the following choices BEST identifies your responsibilities? (Select TWO.)
- A. Hardening systems
 - B. Detecting security incidents
 - C. Preventing incidents
 - D. Implementing monitoring controls
13. A security professional has reported an increase in the number of tailgating violations into a secure data center. What can prevent this?
- A. CCTV
 - B. Mantrap
 - C. Proximity card
 - D. Cipher lock
14. You are redesigning your password policy. You want to ensure that users change their passwords regularly, but they are unable to reuse passwords. What settings should you configure? (Select THREE.)
- A. Maximum password age
 - B. Password length
 - C. Password history
 - D. Password complexity
 - E. Minimum password age
15. An outside security auditor recently completed an in-depth security audit on your network. One of the issues he reported was related to passwords. Specifically, he found the following passwords used on the network: Pa\$\$, 1@W2, and G7bT3. What should be changed to avoid the problem shown

with these passwords?

- A. Password complexity
- B. Password length
- C. Password history
- D. Password reuse

16. A recent security audit discovered several apparently dormant user accounts. Although users could log on to the accounts, no one had logged on to them for more than 60 days. You later discovered that these accounts are for contractors who work approximately one week every quarter. What is the BEST response to this situation?

- A. Remove the account expiration from the accounts.
- B. Delete the accounts.
- C. Reset the accounts.
- D. Disable the accounts.

17. Your organization routinely hires contractors to assist with different projects. Administrators are rarely notified when a project ends and contractors leave. Which of the following is the BEST choice to ensure that contractors cannot log on with their account after they leave?

- A. Enable account expiration.
- B. Enable an account enablement policy.
- C. Enable an account recovery policy.
- D. Enable generic accounts.

18. Developers are planning to develop an application using role-based access control. Which of the following would they MOST likely include in their planning?

- A. A listing of labels reflecting classification levels
- B. A requirements list identifying need to know
- C. A listing of owners
- D. A matrix of functions matched with their required privileges

19. An organization has implemented an access control model that enforces permissions based on data labels assigned at different levels. What type of model is this?

- A. DAC
- B. MAC
- C. Role-BAC
- D. Rule-BAC

20. Your organization's security policy requires that PII data at rest and PII data in transit be encrypted. Of the following choices, what would the organization use to achieve these objectives? (Select TWO.)

- A. FTP
- B. SSH
- C. SMTP
- D. PGP/GPG
- E. HTTP

21. Which of the following list of protocols use TCP port 22 by default?

- A. FTPS, TLS, SCP
- B. SCP, SFTP, FTPS
- C. HTTPS, SSL, TLS
- D. SSH, SCP, SFTP
- E. SCP, SSH, SSL

22. Bart wants to block access to all external web sites. Which port should he block at the firewall?

- A. TCP 22
- B. TCP 53
- C. UDP 69
- D. TCP 80

23. You need to manage a remote server. Which of the following ports should you open on the firewall between your system and the remote server?

- A. 25 and 3389
- B. 22 and 443
- C. 22 and 3389
- D. 21 and 23

24. While reviewing logs on a firewall, you see several requests for the AAAA record of gcgapremium.com. What is the purpose of this request?

- A. To identify the IPv4 address of gcgapremium.com
- B. To identify the IPv6 address of gcgapremium.com
- C. To identify the mail server for gcgapremium.com
- D. To identify any aliases used by gcgapremium.com

25. Your organization has several switches used within the network. You need to implement a security control to secure the switch from physical access. What should you do?

- A. Disable unused ports.
- B. Implement an implicit deny rule.
- C. Disable STP.
- D. Enable SSH.

26. You are configuring a switch and need to ensure that only authorized devices can connect to it and access the network through this switch. Which of the following is the BEST choice to meet this goal?

- A. Implement 802.1x
- B. Use a Layer 3 switch.
- C. Create a VLAN
- D. Enable RSTP.

27. You need to configure a UTM security appliance to restrict access to peer-to-peer file sharing web sites. What are you MOST likely to configure?

- A. Content inspection
- B. Malware inspection
- C. URL filter
- D. Stateless inspection

28. Your organization has implemented a network design that allows internal computers to share one public IP address. Of the following choices, what did they MOST likely implement?

- A. PAT
- B. STP
- C. DNAT
- D. TLS

29. What would you configure on a Layer 3 device to allow FTP traffic to pass through?

- A. Router
- B. Implicit deny
- C. Port security
- D. Access control list

30. What type of device would have the following entries used to define its operation?

permit IP any any eq 80

permit IP any any eq 443

deny IP any any

- A. Layer 2 switch
- B. Proxy server
- C. Web server
- D. Firewall

31. You are preparing to deploy an anomaly-based detection system to monitor network activity. What would you create first?

- A. Flood guards
- B. Signatures
- C. Baseline
- D. Honeypot

32. A security company wants to gather intelligence about current methods attackers are using against its clients. What can it use?

- A. Vulnerability scan
- B. Honeynet
- C. MAC address filtering
- D. Evil twin

33. Lisa oversees and monitors processes at a water treatment plant using SCADA systems. Administrators recently discovered malware on her system that was connecting to the SCADA systems. Although they removed the malware, management is still concerned. Lisa needs to continue using her system and it's not possible to update the SCADA systems. What can mitigate this risk?

- A. Install HIPS on the SCADA systems.
- B. Install a firewall on the border of the SCADA network.
- C. Install a NIPS on the border of the SCADA network.
- D. Install a honeypot on the SCADA network.

34. Your organization maintains a separate wireless network for visitors in a conference room. However, you have recently noticed that people are connecting to this network even when there aren't any visitors in the conference room. You want to prevent these connections, while maintaining easy access for visitors in the conference room. Which of the following is the BEST solution?

- A. Disable SSID broadcasting.
- B. Enable MAC filtering.
- C. Use wireless jamming.
- D. Reduce antenna power.

35. Which of the following represents the BEST action to increase security in a wireless network?

- A. Replace dipole antennas with Yagi antennas.
- B. Replace TKIP with CCMP.
- C. Replace WPA with WEP.
- D. Disable SSID broadcast.

36. Your organization is hosting a wireless network with an 802.1x server using PEAP. On Thursday, users report they can no longer access the wireless network. Administrators verified the network configuration matches the baseline, there aren't any hardware outages, and the wired network is operational. Which of the following is the MOST likely cause for this problem?

- A. The RADIUS server certificate expired.
- B. DNS is providing incorrect host names.
- C. DHCP is issuing duplicate IP addresses.
- D. MAC filtering is enabled.

37. You are planning a wireless network for a business. A core requirement is to ensure that the solution encrypts user credentials when users enter their usernames and passwords. Which of the following BEST meets this requirement?

- A. WPA2-PSK

- B. WEP over PEAP
- C. WPS with LEAP
- D. WPA2 over EAP-TTLS

38. A small business owner modified his wireless router with the following settings:

PERMIT 1A:2B:3C:4D:5E:6F

DENY 6F:5E:4D:3C:2B:1A

After saving the settings, an employee reports that he cannot access the wireless network anymore. What is the MOST likely reason that the employee cannot access the network?

- A. IP address filtering
- B. Hardware address filtering
- C. Port filtering
- D. URL filtering

39. Homer recently implemented a wireless network in his home using WEP. He asks you for advice. Which of the following is the BEST advice you can give him?

- A. He should not use WEP because it uses a weak encryption algorithm.
- B. He should also ensure he disables SSID broadcast for security purposes.
- C. He should ensure it is in Enterprise mode.
- D. He should not use WEP because it implements weak IVs for encryption keys.

40. Which of the following is an attack against a mobile device?

- A. War chalking
- B. SSID hiding
- C. Evil twin
- D. Bluejacking

41. A network administrator needs to open a port on a firewall to support a VPN using PPTP. What ports should the administrator open?

- A. UDP 47
- B. TCP 50
- C. TCP 1723
- D. UDP 1721

42. Attackers recently attacked a web server hosted by your organization. Management has tasked administrators with reducing the attack surface of this server to prevent future attacks. Which of the following will meet this goal?

- A. Disabling unnecessary services
- B. Installing and updating antivirus software
- C. Identifying the baseline
- D. Installing a NIDS

43. Network administrators identified what appears to be malicious traffic coming from an internal computer, but only when no one is logged on to the computer. You suspect the system is infected with malware. It periodically runs an application that attempts to connect to web sites over port 80 with Telnet. After comparing the computer with a list of services from the standard image, you verify this application is very likely the problem. What process allowed you to make this determination?

- A. Banner grabbing
- B. Hardening
- C. Whitelisting
- D. Baselining

44. An updated security policy defines what applications users can install and run on company-issued mobile devices. Which of the following technical controls will enforce this policy?

- A. Whitelisting
- B. Blacklisting
- C. AUP
- D. BYOD

45. You want to test new security controls before deploying them. Which of the following technologies provides the MOST flexibility to meet this goal?

- A. Baselines
- B. Hardening techniques
- C. Virtualization technologies
- D. Patch management programs

46. An organization recently suffered a significant outage after a technician installed an application update on a vital server during peak hours. The server remained down until administrators were able to install a previous version of the application on the server. What could the organization implement to prevent a reoccurrence of this problem?

- A. Do not apply application patches to server applications.
- B. Apply the patches during nonpeak hours.
- C. Apply hardening techniques.
- D. Create a patch management policy.

47. A security analyst is evaluating a critical industrial control system. The analyst wants to ensure the system has security controls to support availability. Which of the following will BEST meet this need?

- A. Using at least two firewalls to create a DMZ
- B. Installing a SCADA system
- C. Implementing control redundancy and diversity
- D. Using an embedded system

48. Of the following choices, what are valid security controls for mobile devices?

- A. Screen locks, device encryption, and remote wipe
- B. Host-based firewalls, pop-up blockers, and SCADA access
- C. Antivirus software, voice encryption, and NAC
- D. Remote lock, NAC, and locking cabinets

49. A new mobile device security policy has authorized the use of employee-owned devices, but mandates additional security controls to protect them if devices are lost or stolen. Which of the following meets this goal?

- A. Screen locks and geo-tagging
- B. Patch management and change management
- C. Screen locks and device encryption
- D. Full device encryption and IaaS

50. You want to deter an attacker from using brute force to gain access to a mobile device. What would you configure?

- A. Remote wiping
- B. Account lockout settings
- C. Geo-tagging
- D. RFID

51. Management within your company is considering allowing users to connect to the corporate network with their personally owned devices. Which of the following represents a security concern with this policy?

- A. Inability to ensure devices are up to date with current system patches
- B. Difficulty in locating lost devices
- C. Cost of the devices

D. Devices might not be compatible with applications within the network

52. Your organization is planning to issue mobile devices to some employees, but they are concerned about protecting the confidentiality of data if the devices are lost or stolen. Which of the following is the BEST way to secure data at rest on a mobile device?

- A. Strong passwords
- B. Hashing
- C. RAID-6
- D. Full device encryption

53. Your organization recently purchased several new laptop computers for employees. You're asked to encrypt the laptop's hard drives without purchasing any additional hardware. What would you use?

- A. TPM
- B. HSM
- C. VM escape
- D. DLP

54. Management within your organization wants to limit documents copied to USB flash drives. Which of the following can be used to meet this goal?

- A. DLP
- B. Content filtering
- C. IPS
- D. Logging

55. Bart installed code designed to enable his account automatically, three days after anyone disables it. What does this describe?

- A. Logic bomb
- B. Rootkit
- C. Armored virus
- D. Ransomware

56. Lisa recently completed an application used by the Personnel department to store PII and other employee information. She programmed in the ability to access this application with a username and password that only she knows, so that she can perform remote maintenance on the application if necessary. What does this describe?

- A. Armored virus
- B. Polymorphic virus
- C. Backdoor
- D. Trojan

57. A recent change in an organization's security policy states that monitors need to be positioned so that they cannot be viewed from outside any windows. What is the purpose of this policy?

- A. Reduce success of phishing
- B. Reduce success of shoulder surfing
- C. Reduce success of dumpster diving
- D. Reduce success of impersonation

58. You are troubleshooting an intermittent connectivity issue with a web server. After examining the logs, you identify repeated connection attempts from various IP addresses. You realize these connection attempts are overloading the server, preventing it from responding to other connections. Which of the following is MOST likely occurring?

- A. DDoS attack
- B. DoS attack
- C. Smurf attack
- D. Salting attack

59. Your organization includes the following statement in the security policy:
“Security controls need to protect against both online and offline password brute force attacks.”
Which of the following controls is the LEAST helpful to meet these goals?
- A. Account expiration
 - B. Account lockout
 - C. Password complexity
 - D. Password length
60. A code review of a web application discovered that the application is not performing boundary checking. What should the web developer add to this application to resolve this issue?
- A. XSRF
 - B. XSS
 - C. Input validation
 - D. Fuzzing
61. A web developer is using methods to validate user input in a web site application. This ensures the application isn’t vulnerable to all of the following attacks except one. Which of the following attacks are NOT prevented by validating user input?
- A. XSS
 - B. SQL injection
 - C. Buffer overflow
 - D. Command injection
 - E. Whaling
62. Checking the logs of a web server, you see the following entry:
198.252.69.129 --[1/Sep/2013:05:20]"GET /index.php?username=////////////////BBBBBBBBBBBBCCCCCHTTP/1.1"
"http://gcgapremium.com/security/" "Chrome31"
- Which of the following is the BEST choice to explain this entry?
- A. A SQL injection attack
 - B. A pharming attack
 - C. A phishing attack
 - D. A buffer overflow attack
63. Looking at logs for an online web application, you see that someone has entered the following phrase into several queries:
' or '1'='1' --
- Which of the following is the MOST likely explanation for this?
- A. A buffer overflow attack
 - B. An XSS attack
 - C. A SQL injection attack
 - D. An LDAP injection attack
64. A security tester is using fuzzing techniques to test a software application. Which of the following does fuzzing use to test the application?
- A. Formatted input
 - B. Unexpected input
 - C. Formatted output
 - D. Unexpected output
65. An organization has purchased fire insurance to manage the risk of a potential fire. What method are they using?
- A. Risk acceptance
 - B. Risk avoidance
 - C. Risk deterrence
 - D. Risk mitigation

E. Risk transference

66. You are asked to identify the number of times a specific type of incident occurs per year. Which of the following BEST identifies this?

- A. ALE
- B. ARO
- C. MTTF
- D. SLE

67. Lisa needs to calculate the total ALE for a group of servers used in the network. During the past two years, five of the servers failed. The hardware cost to replace each server is \$3,500, and the downtime has resulted in \$2,500 of additional losses. What is the ALE?

- A. \$7,000
- B. \$8,000
- C. \$15,000
- D. \$30,000

68. Security experts at your organization have determined that your network has been repeatedly attacked from multiple entities in a foreign country. Research indicates these are coordinated and sophisticated attacks. What BEST describes this activity?

- A. Fuzzing
- B. Sniffing
- C. Spear phishing
- D. Advanced persistent threat

69. Bart is performing a vulnerability assessment. Which of the following BEST represents the goal of this task?

- A. Identify services running on a system.
- B. Determine if vulnerabilities can be exploited.
- C. Determine if input validation is in place.
- D. Identify the system's security posture.

70. You need to ensure that several systems have all appropriate security controls and patches. However, your supervisor specifically told you not to attack or compromise any of these systems. Which of the following is the BEST choice to meet these goals?

- A. Vulnerability scan
- B. Penetration test
- C. Command injection
- D. Virus scan

71. Which of the following tools is the MOST invasive type of testing?

- A. Pentest
- B. Protocol analyzer
- C. Vulnerability scan
- D. Host enumeration

72. A security professional is testing the functionality of an application, but does not have any knowledge about the internal coding of the application. What type of test is this tester performing?

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

73. Testers are analyzing a web application your organization is planning to deploy. They have full access to product documentation, including the code and data structures used by the application. What type of test will they MOST likely perform?

- A. Gray box
- B. White box

- C. Black box
- D. White hat

74. A network administrator is attempting to identify all traffic on an internal network. Which of the following tools is the BEST choice?

- A. Black box test
- B. Protocol analyzer
- C. Penetration test
- D. Baseline review

75. Your organization security policy requires that personnel notify security administrators if an incident occurs. However, this is not occurring consistently. Which of the following could the organization implement to ensure security administrators are notified in a timely manner?

- A. Routine auditing
- B. User rights and permissions reviews
- C. Design review
- D. Incident response team

76. A security administrator is reviewing an organization's security policy and notices that the policy does not define a time frame for reviewing user rights and permissions. Which of the following is the MINIMUM time frame that she should recommend?

- A. At least once a year
- B. At least once every five years
- C. Anytime an employee leaves the organization
- D. Anytime a security incident has been identified

77. Security personnel recently performed a security audit. They identified several employees who had permissions for previously held jobs within the company. What should the organization implement to prevent this in the future?

- A. Role-BAC model
- B. Account disablement policy
- C. Vulnerability assessment
- D. Account management controls

78. You are a technician at a small organization. You need to add fault-tolerance capabilities within the business to increase the availability of data. However, you need to keep costs as low as possible. Which of the following is the BEST choice to meet these needs?

- A. Failover cluster
- B. RAID-6
- C. Backups
- D. UPS

79. An organization needs to identify a continuity of operations plan that will allow it to provide temporary IT support during a disaster. The organization does not want to have a dedicated site. Which of the following provides the best solution?

- A. Cold site
- B. Warm site
- C. Hot site
- D. Mobile site

80. Monty Burns is the CEO of the Springfield Nuclear Power Plant. What would the company have in place in case something happens to him?

- A. Business continuity planning
- B. Succession planning
- C. Separation of duties
- D. IT contingency planning

81. A continuity of operations plan for an organization includes the use of a warm site. The BCP coordinator wants to verify that the organization's

backup data center is prepared to implement the warm site if necessary. Which of the following is the BEST choice to meet this need?

- A. Perform a review of the disaster recovery plan.
- B. Ask the managers of the backup data center.
- C. Perform a disaster recovery exercise.
- D. Perform a test restore.

82. Users are complaining of intermittent connectivity issues. When you investigate, you discover that new network cables for these user systems were run along several fluorescent lights. What environmental control will resolve this issue?

- A. HVAC system
- B. Fire suppression
- C. Humidity controls
- D. EMI shielding

83. A software company occasionally provides application updates and patches via its web site. It also provides a checksum for each update and patch. Which of the following BEST describes the purpose of the checksum?

- A. Availability of updates and patches
- B. Integrity of updates and patches
- C. Confidentiality of updates and patches
- D. Integrity of the application

84. A function converts data into a string of characters and the string of characters cannot be reversed to re-create the original data. What type of function is this?

- A. Symmetric encryption
- B. Asymmetric encryption
- C. Stream cipher
- D. Hashing

85. Which of the following is a symmetric encryption algorithm that encrypts data one bit at a time?

- A. Block cipher
- B. Stream cipher
- C. AES
- D. DES
- E. MD5

86. A supply company has several legacy systems connected together within a warehouse. An external security audit discovered the company is using DES and mandated the company upgrade DES to meet minimum security requirements. The company plans to replace the legacy systems next year, but needs to meet the requirements from the audit. Which of the following is MOST likely to be the simplest upgrade for these systems?

- A. AES
- B. HMAC
- C. 3DES
- D. SSL

87. Network administrators in your organization need to administer firewalls, security appliances, and other network devices. These devices are protected with strong passwords, and the passwords are stored in a file listing these passwords. Which of the following is the BEST choice to protect this password list?

- A. File encryption
- B. Database field encryption
- C. Full database encryption
- D. Whole disk encryption

88. Bart, an employee at your organization, is suspected of leaking data to a competitor. Investigations indicate he sent several email messages

containing pictures of his dog. Investigators have not been able to identify any other suspicious activity. Which of the following is MOST likely occurring?

- A. Bart is copying the data to a USB drive.
- B. Bart is encrypting the data.
- C. Bart is leaking data using steganography.
- D. Bart is sending the data as text in the emails.

89. You are planning to encrypt data in transit with IPsec. Which of the following is MOST likely to be used with IPsec?

- A. HMAC
- B. Blowfish
- C. Twofish
- D. MD5

90. Bart wants to send a secure email to Lisa, so he decides to encrypt it. He wants to ensure that only Lisa can decrypt it. Which of the following does Lisa need to meet this requirement?

- A. Bart's public key
- B. Bart's private key
- C. Lisa's public key
- D. Lisa's private key

91. An organization requested bids for a contract and asked companies to submit their bids via email. After winning the bid, Acme realized it couldn't meet the requirements of the contract. Acme instead stated that it never submitted the bid. Which of the following would provide proof to the organization that Acme did submit the bid?

- A. Digital signature
- B. Integrity
- C. Repudiation
- D. Encryption

92. Application developers are creating an application that requires users to log on with strong passwords. The developers want to store the passwords in such a way that it will thwart brute force attacks. Which of the following is the BEST solution?

- A. 3DES
- B. MD5
- C. PBKDF2
- D. Database fields

93. A web site is using a certificate. Users have recently been receiving errors from the web site indicating that the web site's certificate is revoked. Which of the following includes a list of certificates that have been revoked?

- A. CRL
- B. CA
- C. OCSP
- D. CSR

94. Which of the following is a management control?

- A. Encryption
- B. Security policy
- C. Least privilege
- D. Change management

95. Security personnel recently identified potential fraud committed by a network administrator. Investigators discovered this administrator performs several job functions within the organization, including database administration and application development. Which of the following is the BEST solution to reduce risk associated with this activity?

- A. Mandatory vacations

- B. Mandatory access control
- C. Change management
- D. Separation of duties

96. Security experts want to reduce risks associated with updating critical operating systems. Which of the following will BEST meet this goal?

- A. Load balancing
- B. Change management
- C. Incident management
- D. Key management

97. Your company is considering implementing SSO capabilities to company applications and linking them to a social media site. When implemented, users can log on to Facebook and then access company applications without logging on again. What is a potential risk related to this plan?

- A. A data breach exposing passwords on the company site will affect the social media site.
- B. SAML lacks adequate security when used on the Internet.
- C. XML lacks adequate security when used on the Internet.
- D. A data breach exposing passwords on the social media site will affect the company application.

98. You work as a help-desk professional in a large organization. You have begun to receive an extraordinary number of calls from employees related to malware. Using common incident response procedures, what should be your FIRST response?

- A. Preparation
- B. Identification
- C. Escalation
- D. Mitigation

99. A technician confiscated an employee's computer after management learned the employee had unauthorized material on his system. Later, a security expert captured a forensic image of the system disk. However, the security expert reported the computer was left unattended for several hours before he captured the image. Which of the following is a potential issue if this incident goes to court?

- A. Chain of custody
- B. Order of volatility
- C. Time offset
- D. Lack of metrics

100. Social engineers have launched several successful phone-based attacks against your organization resulting in several data leaks. Which of the following would be the MOST effective at reducing the success of these attacks?

- A. Implement a BYOD policy.
- B. Update the AUP.
- C. Provide training on data handling.
- D. Implement a program to increase security awareness.

Assessment Exam Answers

When checking your answers, take the time to read the explanations. Understanding the explanations will help ensure you're prepared for the live exam. The explanation also shows the chapter or chapters where you can get more detailed information on the topic.

1. **C.** The administrator should ensure systems are not susceptible to unauthorized changes, an element of integrity. A security program should address the three core security principles of confidentiality, integrity, and availability; the system in the example is already addressing confidentiality and availability. Protecting data and securing data to prevent unauthorized disclosure addresses confidentiality. Ensuring critical systems provide uninterrupted service addresses availability. See Chapter 1.
2. **B.** You can maintain confidentiality of any data, including Personally Identifiable Information (PII) with encryption. Hashes provide integrity, not confidentiality. A digital signature provides authentication, non-repudiation, and integrity. A redundant array of inexpensive disks (RAID) provides higher availability for a disk subsystem. See Chapters 1 and 10.
3. **B.** She is pursuing integrity by verifying the configuration files have not changed. By verifying that the hashes are the same, she also verifies that the configuration files are the same. Confidentiality is enforced with encryption, access controls, and steganography. Availability ensures systems are up and operational when needed. Safety goals help ensure the safety of personnel and/or other assets. See Chapters 1 and 10.
4. **B.** Layered security (or defense in depth) implements multiple controls to provide several layers of protection. In this case, the antivirus software provides one layer of protection while the firewall and the intrusion detection system (IDS) provide additional layers. Implicit deny blocks access unless it has been explicitly allowed. Least privilege ensures that users are granted only the access they need to perform their jobs, and no more. A flood guard attempts to block SYN Flood attacks. See Chapter 1.
5. **D.** In this scenario, it's best to create a temporary password that expires after first use, which forces the user to create a new password. It's not necessary to verify the user's account exists, but the help-desk professional should verify the identity of the user. Passwords should not be available in such a way that allows help-desk professionals to look them up. It is not necessary to disable a user account to reset the password. See Chapter 1.
6. **D.** A hardware token (such as an RSA token) uses a one-time password for authentication in the something you have factor of authentication. Biometric methods are in the something you are factor of authentication, such as a fingerprint. A PIN and a password are both in the something you know factor of authentication and do not require a hardware token. See Chapter 1.
7. **C.** A retina scan is a biometric method of authentication in the something you are factor of authentication. You need to combine two or more factors of authentication for dual-factor and multifactor authentication. A Time-based One-Time Password (TOTP) is a protocol used to create passwords that expire after 30 seconds. See Chapter 1.
8. **B.** Users authenticate with two factors of authentication in this scenario, which is multifactor authentication or dual-factor authentication. The smart card is in the something you have factor of authentication, and the PIN is in the something you know factor of authentication. They are using more than a single factor. Mutual authentication is when both entities in the authentication process authenticate with each other, but it doesn't apply in this situation. A Time-based One-Time Password (TOTP) is a protocol used to create passwords that expire after 30 seconds. See Chapter 1.
9. **B.** Remote Authentication Dial-In User Service (RADIUS) is an authentication, authorization, and accounting (AAA) protocol and is the best choice. TACACS+ is proprietary to Cisco, so it won't support different vendor solutions. Diameter is preferable to RADIUS, but there is no such thing as a Circumference protocol. SAML is an SSO solution used with web-based applications. See Chapter 1.
10. **C.** This describes a single sign-on (SSO) solution in which users only have to log on once. Same sign-on indicates users can access multiple systems using the same credentials, but they still have to enter their credentials again each time they access a new resource. Security Assertion Markup Language (SAML) is an SSO solution used for web-based applications, but not all SSO solutions are using SAML. Biometrics is a method of authentication, such as a fingerprint, but it isn't an SSO solution. See Chapter 1.
11. **C.** Disk encryption is a strong technical control that can mitigate potential data losses if mobile devices are lost or stolen. Cable locks are preventive controls that can prevent the theft of mobile devices such as laptops, but they don't protect the data after the device is stolen. A risk assessment is a management control. Hardening systems helps make them more secure than their default configuration, but doesn't necessarily protect data after the device is lost. See Chapters 2 and 5.
12. **B, D.** Monitoring security logs and analyzing trend reports are detective controls with the goal of detecting security incidents. Installing closed-circuit television (CCTV) systems is one example of implementing a monitoring control. Hardening a system is a preventive control that includes several steps such as disabling unnecessary services, but the scenario doesn't describe these steps. Preventive controls attempt to prevent incidents, but the scenario describes detective controls. See Chapter 2.
13. **B.** A mantrap is highly effective at preventing unauthorized entry and can also be used to prevent tailgating. CCTV provides video surveillance and it can record unauthorized entry, but it can't prevent it. A proximity card is useful as an access control mechanism, but it won't prevent tailgating, so it isn't as useful as a mantrap. A cipher lock is a door access control, but it can't prevent tailgating. See Chapter 2.
14. **A, C, E.** The maximum password age ensures users change their passwords regularly. The password history records previously used passwords

- (such as the last 24 passwords) to prevent users from reusing the same passwords. The minimum password age prevents users from changing their password repeatedly to get back to their original password and should be used with the password history setting. Password length requires a minimum number of characters in a password. Password complexity requires a mix of uppercase and lowercase letters, numbers, and special characters. See Chapter 2.
15. **B.** The password policy should be changed to increase the minimum password length of passwords. These passwords are only four and five characters long, which is too short to provide adequate security. They are complex because they include a mixture of at least three of the following character types: uppercase letters, lowercase letters, numbers, and special characters. Password history and password reuse should be addressed if users are reusing the same passwords, but the scenario doesn't indicate this is a problem. See Chapter 2.
16. **D.** The best response is to disable the accounts and then enable them when needed by the contractors. Ideally, the accounts would include an expiration date so that they would automatically expire when no longer needed, but the scenario doesn't indicate the accounts have an expiration date. Because the contractors need to access the accounts periodically, it's better to disable them rather than deleting them. Reset the accounts implies you are changing the password, but this isn't needed. See Chapter 2.
17. **A.** The best choice is to enable account expiration so that the contractor accounts are automatically disabled at the end of their projected contract time period. If contracts are extended, it's easy to enable the account and reset the account expiration date. Account disablement policies help ensure that any user accounts (not just contractors) are disabled when the user leaves the organization, but an account enablement policy isn't a valid term. An account recovery policy allows administrators to recover accounts and associated security keys for ex-employees. It's best to prohibit the use of generic accounts (such as the Guest account), so enabling generic accounts is not recommended. See Chapter 2.
18. **D.** A matrix of functions, roles, or job titles matched with the required access privileges for each of the functions, roles, or job titles is a common planning document for a role-based access control model. The mandatory access control (MAC) model uses sensitivity labels and classification levels. MAC is effective at restricting access based on a need to know. The discretionary access control model specifies that every object has an owner and it might identify owners in a list. See Chapter 2.
19. **B.** The mandatory access control (MAC) model uses labels assigned at different levels to restrict access. The discretionary access control (DAC) model assigns permissions based on object ownership. The role-based access control (role-BAC) model uses group-based privileges. The rule-based access control (rule-BAC) model uses rules that trigger in response to events. See Chapter 2.
20. **B, D.** You can use Secure Shell (SSH) to encrypt Personally Identifiable Information (PII) data when transmitting it over the network (data in transit). While Pretty Good Privacy (PGP)/GNU Privacy Guard (GPG) is primarily used to encrypt email, it can also be used to encrypt data at rest. File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), and Hypertext Transfer Protocol (HTTP) transmit data in cleartext unless they are combined with an encryption protocol. See Chapters 3 and 10.
21. **D.** Secure Shell (SSH) uses Transmission Control Protocol (TCP) port 22 by default. Secure Copy (SCP) and Secure File Transfer Protocol (SFTP) both use SSH for encryption so they also use port 22 by default. File Transfer Protocol Secure (FTPS) uses either Secure Sockets Layer (SSL) or Transport Layer Security (TLS), typically on ports 989 or 990. Hypertext Transfer Protocol Secure (HTTPS) uses SSL or TLS on port 443. TLS and SSL do not have a default port by themselves, but instead use a default port based on the protocols they are encrypting.
22. **D.** He should block port 80 because web sites use Hypertext Transfer Protocol (HTTP) over TCP port 80. Secure Shell (SSH) uses TCP port 22. Domain Name System (DNS) uses TCP port 53 for zone transfers. Trivial File Transfer Protocol (TFTP) uses UDP port 69.
23. **C.** You can manage a remote server using Secure Shell (SSH) on TCP port 22 and Remote Desktop Protocol (RDP) on TCP port 3389. You could also use Telnet on TCP port 23, but SSH is the preferred alternative. Simple Mail Transfer Protocol (SMTP) uses TCP port 25. Hypertext Transfer Protocol Secure (HTTPS) uses TCP port 443. File Transfer Protocol (FTP) uses TCP port 21.
24. **B.** A Domain Name System (DNS) AAAA record identifies the IPv6 address of a given name. An A record identifies the IPv4 address of a given name. An MX record identifies a mail server. A CNAME record identifies aliases.
25. **A.** You can provide added security by disabling unused physical ports on the switch. If someone gains physical access to the switch by plugging in a computer to one of its unused ports, that person will not be able to connect to the network. An implicit deny rule is placed at the end of an access control list on a router to deny traffic that hasn't been explicitly allowed, but it doesn't not affect physical ports differently. Spanning Tree Protocol (STP) prevents switching loop problems and should be enabled. Secure Shell (SSH) encrypts traffic but doesn't protect a switch.
26. **A.** An 802.1x server provides port-based authentication and can prevent unauthorized devices from connecting to a network. Although you can configure an 802.1x server with a VLAN to redirect unauthorized clients, the VLAN by itself will not block unauthorized devices. A Layer 3 switch does not provide port-based authentication. Rapid Spanning Tree Protocol (RSTP) will prevent switching loop problems but doesn't authenticate clients.
27. **C.** You would most likely configure the Uniform Resource Locator (URL) filter on the unified thread management (UTM) security appliance. This would block access to the peer-to-peer sites based on their URL. Content inspection and malware inspection focus on inspecting the data as it passes through the UTM, but they do not block access to sites. Stateless inspection is packet filtering and would be extremely difficult to configure on a firewall for all peer-to-peer web sites.
28. **A.** Port Address Translation (PAT) is a form of Network Address Translation (NAT) and it allows many internal devices to share one public IP address. Dynamic Network Address Translation (DNAT) uses multiple public IP addresses instead of just one. Spanning Tree Protocol (STP) prevents

switch loop problems and is unrelated to sharing IPs. Transport Layer Security (TLS) secures transmissions for data in transit.

29. **D.** You would configure an access control list (ACL) to allow traffic in or out of a network. A router is a Layer 3 device and you would configure the ACL on the router. The last rule in the ACL would be implicit deny to block all other traffic. Port security protects ports by disabling unused ports or using 802.1x, but it cannot block specific types of traffic.

30. **D.** These are rules in an access control list (ACL) for a firewall. The first two rules indicate that traffic from any IP address, to any IP address, using ports 80 or 443 is permitted or allowed. The final rule is also known as an implicit deny rule and is placed last in the ACL. It ensures that all traffic that hasn't been previously allowed is denied. Layer 2 switches do not use ACLs. A proxy server would not use an ACL, although it would use ports 80 and 443 for Hypertext Transfer Protocol (HTTP) and HTTP Secure (HTTPS), respectively. A web server wouldn't use an ACL, although it would also use ports 80 and 443. See Chapter 8.

31. **C.** An anomaly-based (also called heuristic or behavior-based) detection system compares current activity with a previously created baseline to detect any anomalies or changes. Flood guards help protect against SYN flood attacks. Signature-based systems use signatures similar to antivirus software. A honeypot is a server designed to look valuable to an attacker and can divert attacks. See Chapter 4.

32. **B.** A honeynet is a fake network designed to look valuable to attackers and can help security personnel learn about current attack methods. In this scenario, the security company can install honeynets in its customers' networks to lure the attackers. A vulnerability scan detects vulnerabilities, but attackers may not try to exploit them. Media access control (MAC) address filtering is a form of network access control, but can't be used to detect or learn about attacks. An evil twin is a rogue access point with the same SSID as an authorized access point. See Chapter 4.

33. **C.** A network intrusion prevention system (NIPS) installed on the supervisory control and data acquisition (SCADA) network can intercept malicious traffic coming into the network and is the best choice of those given. The scenario states you cannot update the SCADA systems, so you cannot install a host-based IPS (HIPS) on any of them. A firewall provides a level of protection. However, it wouldn't be able to differentiate between valid traffic sent by Lisa and malicious traffic sent by malware from Lisa's system. A honeypot might be useful to observe malicious traffic, but wouldn't prevent it. See Chapter 4.

34. **D.** Reducing the antenna power will make it more difficult for users outside of the conference room to connect, but will not affect visitors in the conference room. Disabling service set identifier (SSID) broadcasting will require visitors to know the SSID and enter it in their device, making it more difficult to access the wireless network. Enabling media access control (MAC) address filtering will block visitors until an administrator adds their MAC address. Wireless jamming will prevent all mobile devices from connecting to the wireless network. See Chapter 4.

35. **B.** Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP) provides stronger encryption than Temporal Key Integrity Protocol (TKIP) and is the best choice. Replacing omnidirectional dipole antennas with directional Yagi antennas doesn't necessarily increase security and will likely limit availability. Wired Equivalent Privacy (WEP) should not be used and is not an improvement over Wi-Fi Protected Access (WPA). Disabling service set identifier (SSID) broadcast hides the network from casual users, but is not a security step. See Chapter 4.

36. **A.** The most likely cause is that the Remote Authentication Dial-In User Service (RADIUS) server certificate expired. An 802.1x server is implemented as a RADIUS server and Protected Extensible Authentication Protocol (PEAP) requires a certificate. If Domain Name System (DNS) or Dynamic Host Configuration Protocol (DHCP) failed, it would affect both wired and wireless users. Media access control (MAC) address filtering might cause this symptom if all MAC addresses were blocked, but the scenario states that there weren't any network configuration changes. See Chapter 4.

37. **D.** Wi-Fi Protected Access II (WPA2) over Extensible Authentication Protocol (EAP)-Tunneled Transport Layer Security (EAP-TTLS) is the best solution from the available answers. Because users must enter their usernames and passwords, an 802.1x solution is required and EAP-TTLS meets this requirement. WPA2-preshared key (PSK) does not authenticate users based on their usernames. Wired Equivalent Privacy (WEP) is not recommended for use even with Protected EAP (PEAP). Wi-Fi Protected Setup (WPS) is a standard designed to simplify the setup of a wireless network, but it does not implement usernames, and Cisco recommends using stronger protocols rather than Lightweight EAP (LEAP). See Chapter 4.

38. **B.** Media access control (MAC) address filtering can block or allow access based on a device's MAC address, also known as the hardware address. Both addresses in the scenario are MAC addresses. These addresses are not Internet Protocol (IP) addresses, port numbers, or Uniform Resource Locators (URLs). See Chapter 4.

39. **D.** Wired Equivalent Privacy (WEP) is not recommended for use and one of the reasons is due to weak initialization vectors (IVs) used for key transmission. It uses the RC4 stream cipher, which is a strong encryption algorithm. Disabling the service set identifier (SSID) broadcast will hide the network from casual users, but it does not provide additional security. WEP doesn't support Enterprise mode. See Chapter 4.

40. **D.** Bluejacking is the practice of sending unsolicited messages to other Bluetooth devices. War chalking is the practice of marking the location of wireless networks, sometimes using chalk. You can disable service set identifier (SSID) broadcasting to hide the SSID from casual users, but this isn't an attack. An evil twin is a rogue access point with the same SSID as a legitimate access point. It can be used to launch attacks against any wireless devices, but it isn't an attack against only mobile devices. See Chapter 4.

41. **C.** A virtual private network (VPN) using Point-to-Point Tunneling Protocol (PPTP) requires Transmission Control Protocol (TCP) port 1723 open. It would also need protocol ID 47 open, but the protocol ID is not a port. Internet Protocol security (IPsec) uses protocol ID 50. See Chapters 3 and 4.

42. **A.** Disabling unnecessary services is a primary method of reducing the attack surface of a host. Installing up-to-date antivirus software is valid

preventive control, but it doesn't reduce the attack surface. Identifying the baseline should be done after disabling unnecessary services. A network-based intrusion detection system (NIDS) helps protect the server, but it doesn't reduce its attack surface. See Chapter 5.

43. **D.** The standard image is the baseline and by comparing the list of services in the baseline with the services running on the suspect computer, you can identify unauthorized services. In this scenario, Telnet must not be in the baseline, but it is running on the suspect computer. It's possible an attacker has hijacked the computer to perform banner-grabbing attacks against external web sites, but banner grabbing doesn't verify the problem on the computer. Hardening makes a computer more secure than the default configuration, but it is done before creating a baseline. Whitelisting identifies authorized applications and prevents unauthorized applications from running. See Chapter 5.

44. **A.** Whitelisting identifies authorized software and prevents users from installing or running any other software. Blacklisting identifies what isn't authorized, but in this scenario the policy defines what can be installed, not what cannot be installed. An acceptable use policy is not a technical control. Bring your own device (BYOD) doesn't apply here because the devices are company-issued. See Chapter 5.

45. **C.** Virtualization provides a high degree of flexibility when testing security controls because testers can easily rebuild virtual systems or revert them using a snapshot. Baselines provide a known starting point, but aren't flexible because they stay the same. Hardening techniques make systems more secure than their default configuration. Patch management programs ensure patches are deployed, but do not test security controls. See Chapter 5.

46. **D.** An application patch management policy includes plans for identifying, testing, scheduling, and deploying updates. Patches are often applied to test systems before they are applied to live production systems and this would prevent this outage. Server applications should be kept up to date with patches. Although applying patches during nonpeak hours is a good recommendation, it would have still caused an outage in this scenario. Hardening techniques makes a system more secure, but won't protect systems from a faulty patch. See Chapter 5.

47. **C.** A critical industrial control system implies a supervisory control and data acquisition (SCADA) system and ensuring that the system incorporates diversity into a redundant design will best meet this need of the available choices. A demilitarized zone (DMZ) provides some protection against Internet attacks, but critical industrial control systems rarely have direct Internet access. The goal in the question is to protect the SCADA system, but the SCADA system isn't a security control. The scenario is describing an embedded system. See Chapter 5.

48. **A.** Screen locks, device encryption, and remote wipe are all valid security controls for mobile devices. It's rare for mobile devices to have firewalls, but granting them access to supervisory control and data acquisition (SCADA) systems doesn't protect mobile devices or SCADA systems. Network access control (NAC) provides protection for networks, not mobile devices. See Chapter 5.

49. **C.** Screen locks provide protection for lost devices by making it more difficult for someone to access the device. Device encryption protects the data. Geo-tagging includes location information on pictures posted to social media sites. Patch management keeps devices up to date and change management helps prevent outages from unauthorized changes. Infrastructure as a Service (IaaS) is a cloud computing option. See Chapter 5.

50. **B.** Account lockout settings are useful on any type of device, including mobile devices and desktop systems. An account lockout setting locks a device after a specified number of incorrect password or PIN guesses; some devices can be configured to erase all the data on the device after too many incorrect guesses. Remote wiping erases all the data. Geo-tagging provides geographic location for pictures posted to social media sites. Radio-frequency identification (RFID) can be used for automated inventory control to detect movement of devices. See Chapters 1 and 5.

51. **A.** A core security concern with bring your own device (BYOD) policies is ensuring that they are up to date with current patches and have up-to-date antivirus signature files. Tools are available to locate lost devices even if they are employee-owned. The cost of the devices is not a security concern and not a concern to the company because employees pay for their own devices. Although ensuring that the devices are compatible with network applications is a concern, it only affects availability of the application for a single user. See Chapter 5.

52. **D.** Encryption is the best way to protect data, and full device encryption protects data stored on a mobile device. Although strong passwords are useful, if a thief gets a mobile device, it's just a matter of time before the thief bypasses the password. Hashing is used for integrity, but the confidentiality of the data needs to be protected with encryption. Redundant array of inexpensive disks 6 (RAID-6) can increase availability, but not confidentiality. See Chapter 5.

53. **A.** A Trusted Platform Module (TPM) is included in many new laptops and it provides a mechanism for vendors to perform hard drive encryption. Because the TPM components are included, this solution does not require purchasing additional hardware. An HSM is a removable hardware device and is not included with laptops, so it requires an additional purchase. A VM escape attack runs on a virtual system, and if successful, it allows the attacker to control the physical host server and all other virtual servers on the physical server. A network-based data loss prevention (DLP) system can examine and analyze network traffic and detect if confidential company data is included. See Chapter 5.

54. **A.** A data loss prevention (DLP) solution can limit documents copied to a USB drive using content filters. Many devices, such as unified threat management (UTM) devices use content filters, so content filtering alone won't limit copies sent to a flash drive. An intrusion prevention system (IPS) scans traffic coming into a network to block attacks. Logging can record what documents were copied, but it won't limit copying. See Chapter 5.

55. **A.** A logic bomb is code that executes in response to an event. In this scenario, the logic bomb executes when it discovers the account is disabled (indicating Bart is no longer employed at the company). In this scenario, the logic bomb is creating a backdoor. A rootkit includes hidden processes, but it does not activate in response to an event. An armored virus uses techniques to resist reverse engineering. Ransomware demands payment as ransom. See Chapter 6.

56. **C.** A backdoor provides someone an alternative way of accessing the system, which is exactly what Lisa created in this scenario. It might seem as

though she's doing so with good intentions, but if attackers discover a backdoor, they can exploit it. A virus tries to replicate itself, but this account doesn't have a replication mechanism. A Trojan looks beneficial but includes a malicious component. See Chapter 6.

57. **B.** Shoulder surfing is the practice of viewing data by looking over someone's shoulder and it includes looking at computer monitors. Positioning monitors so that they cannot be viewed through a window reduces this threat. Phishing is an email attack. Dumpster diving is the practice of looking through dumpsters. Social engineers often try to impersonate others to trick them. See Chapter 6.

58. **A.** A distributed denial-of-service (DDoS) attack includes attacks from multiple systems with the goal of depleting the target's resources and this scenario indicates multiple connection attempts from different IP addresses. A DoS attack comes from a single system, and a SYN flood is an example of a DoS attack. A smurf attack doesn't attempt to connect to systems but instead sends pings. Salting is a method used to prevent brute force attacks to discover passwords. See Chapter 7.

59. **A.** Account expiration is not an effective defense against brute force attacks. Account lockout helps protect against online brute force attacks. Password complexity and password length help protect against offline brute force attacks. See Chapters 1 and 7.

60. **C.** The lack of input validation is a common coding error and it includes boundary or limit checking to validate data before using it. Proper input validation prevents many problems such as cross-site request forgery (XSRF), cross-site scripting (XSS), buffer overflow, and command injection attacks. Fuzzing injects extra data and tests the effectiveness of input validation. See Chapter 7.

61. **E.** Whaling is a phishing attack using email that targets executives and cannot be prevented with input validation. Input validation can prevent cross-site scripting (XSS), SQL injection, buffer overflow, and command injection attacks. See Chapter 7.

62. **D.** A buffer overflow attack sends more data or unexpected data to a system in the hopes of overloading it and causing a problem. In this case, it is sending a series of letters as the username (?username=ZZZZ....), which is likely longer than any expected username. Input validation can prevent this from succeeding. A SQL injection attack uses specific SQL code, not random letters or characters. A pharming attack attempts to redirect users from one web site to another web site. A phishing attack sends unwanted email to users. See Chapter 7.

63. **C.** Attackers use the phrase in SQL injection attacks to query or modify databases. A buffer overflow attack sends more data or unexpected data to an application with the goal of accessing system memory. A cross-site scripting (XSS) attack attempts to insert HTML or JavaScript code into a web site or email. A Lightweight Directory Application Protocol (LDAP) injection attack attempts to inject LDAP commands to query a directory service database. See Chapter 7.

64. **B.** Fuzzing sends random or unexpected input into an application to test the application's ability to handle it. Command injection attacks use formatted input. Fuzzing does not test the application using any outputs. See Chapter 7.

65. **E.** Purchasing insurance is a common method of risk transference. Organizations often accept a risk when the cost of the control exceeds the cost of the risk, and the risk that remains is residual risk. An organization can avoid a risk by not providing a service or not participating in a risky activity. Risk deterrence attempts to discourage attacks with preventive controls such as a security guard. Risk mitigation reduces risks through internal controls. See Chapter 8.

66. **B.** The annual rate of occurrence (ARO) is the best choice to identify how many times a specific type of incident occurs in a year. Annual loss expectancy (ALE) identifies the expected monetary loss for an incident and single loss expectancy (SLE) identifies the expected monetary loss for a single incident. $ALE = SLE \times ARO$ and if you know any two of these values, you can identify the third value. For example, $ARO = ALE / SLE$. Mean time to failure (MTTF) is not an annual figure. See Chapter 8.

67. **C.** The annual loss expectancy (ALE) is \$15,000. The single loss expectancy (SLE) is \$6,000 (\$3,500 + \$2,500). The annual rate of occurrence (ARO) is 2.5 (five failures in two years or $5 / 2$). You calculate the ALE as $SLE \times ARO$ ($\$6,000 \times 2.5$). See Chapter 8.

68. **D.** An advanced persistent threat is a group of highly organized individuals, typically from a foreign country, with the ability to coordinate sophisticated attacks. Fuzzing is the practice of sending unexpected input to an application for testing and can be used in a security assessment. Sniffing is the practice of capturing traffic with a protocol analyzer. Spear phishing is a targeted phishing attack. See Chapter 8.

69. **D.** A vulnerability assessment identifies a system or network's security posture. A port scanner identifies services running on a system. A penetration test determines if vulnerabilities can be exploited. Although a vulnerability assessment might verify if input validation methods are in place, it includes much more. See Chapter 8.

70. **A.** A vulnerability scan tests systems and can identify unapplied security controls and patches without attacking or compromising the systems. A penetration test potentially attacks or compromises a system. A command injection attack can also potentially cause damage. A virus scan detects viruses, but it doesn't check for security controls or patches. See Chapter 8.

71. **A.** A pentest (or penetration test) is the invasive type of test listed, and can potentially compromise a system. A protocol analyzer is not invasive, but it cannot determine if security controls are in place. A vulnerability scan can verify if security controls are in place and it does not try to exploit these controls using any invasive methods. Host enumeration identifies hosts on a network, but does not check for security controls. See Chapter 8.

72. **B.** A black box tester does not have prior knowledge when testing an application or network. White box testers have full knowledge and gray box testers have some knowledge. Black hat refers to a malicious attacker. See Chapter 8.

73. **B.** White box testers are provided full knowledge about the product or network they are testing. A black box tester does not have access to product documentation, and a gray box tester would have some access to product documentation. White hat refers to a security professional working within the

law. See Chapter 8.

74. **B.** You can use a protocol analyzer (or sniffer) to capture traffic on a network, and then analyze the capture to identify and quantify all the traffic on the network. Penetration tests (including black box tests) attempt to identify and exploit vulnerabilities. A baseline review can identify changes from standard configurations, but they don't necessarily identify all traffic on a network. See Chapter 8.

75. **A.** Routine auditing of the help desk or administrator logs can discover incidents and then match them with reported incidents. A review of user rights and permissions helps ensure they are assigned and maintained appropriately, but do not help with ensuring incidents are reported correctly. A design review ensures that systems and software are developed properly. An incident response team responds to incidents, but they wouldn't necessarily ensure administrators are informed of incidents. See Chapter 8.

76. **A.** User rights and permissions reviews should occur at least once year, and some organizations do them more often. Every five years is too long. Organizations with a high turnover rate might have employees leaving every week and it's not feasible to do a review that often. Performing a review in a response to incidents won't necessarily prevent incidents. See Chapter 8.

77. **D.** Account management controls ensure that accounts only have the permissions they need and no more, and would ensure that user permissions are removed when users no longer need them. User rights and permission reviews also help ensure the controls are effective. A role-based access control (role-BAC) model uses group-based permissions, but it doesn't force administrators to take a user out of a security group when the user moves to a different job. An account disablement policy ensures accounts are disabled when an employee leaves. A vulnerability assessment might detect this as it reviews the organization's security posture, but it won't prevent it. See Chapters 2 and 8.

78. **B.** A redundant array of inexpensive disks 6 (RAID-6) subsystem provides fault tolerance for disks, and increases data availability. A failover cluster provides fault tolerance for servers and can increase data availability but is significantly more expensive than a RAID subsystem. Backups help ensure data availability, but they do not help with fault tolerance. An uninterruptible power supply (UPS) provides fault tolerance for power, but not necessarily for data. See Chapter 9.

79. **D.** A mobile site is a self-contained transportable unit that can be moved around without having a dedicated site. Cold sites, warm sites, and hot sites are dedicated locations. See Chapter 9.

80. **B.** Succession planning identifies people within an organization who can fill leadership positions if they become vacant. It is also helpful during a disaster by ensuring people understand their roles and responsibilities. A succession planning chart is often in a business continuity plan (BCP), but business continuity planning is much broader than just succession planning. A separation of duties policy separates individual tasks of an overall function between different people. IT contingency planning focuses on recovery of IT systems. See Chapter 9.

81. **C.** The best way to test elements of a business continuity plan (BCP) or disaster recovery plan (DRP) is to test the plan by performing a disaster recovery exercise. Asking managers if they are ready and reviewing the plan are both helpful, but not as effective as an exercise. Performing a test restore verifies the backup capabilities, but not necessarily the steps required when implementing a warm site. See Chapter 9.

82. **D.** Electromagnetic interference (EMI) shielding provides protection against EMI sources such as fluorescent lights. Heating, ventilation, and air conditioning systems provide protection from overheating. Fire suppression systems provide protection from fire. Humidity controls provide protection against electrostatic discharge (ESD) and condensation. See Chapter 9.

83. **B.** The checksum (also known as a hash) provides integrity for the patches and updates so that users can verify they have not been modified. Installing patches and updates increases the availability of the application. Confidentiality is provided by encryption. The checksums are for the updates and patches, so they do not provide integrity for the application. See Chapter 10.

84. **D.** A hash function creates a string of characters (typically displayed in hexadecimal) when executed against a file or message, and hashing functions cannot be reversed to re-create the original data. Encryption algorithms (including symmetric encryption, asymmetric encryption, and stream ciphers) create ciphertext from plaintext data, but they include decryption algorithms to re-create the original data. See Chapter 10.

85. **B.** A stream cipher encrypts data a single bit or a single byte at a time and is more efficient when the size of the data is unknown, such as streaming audio or video. A block cipher encrypts data in specific-sized blocks, such as 64-bit blocks or 128-bit blocks. Advanced Encryption Standard (AES) and Data Encryption Standard (DES) are block ciphers. Message Digest 5 (MD5) is a hashing algorithm. See Chapter 10.

86. **C.** The simplest upgrade is Triple Data Encryption Standard (3DES). Advanced Encryption Standard (AES) is stronger, but considering these are legacy systems, their hardware is unlikely to support AES and 3DES is a suitable alternative. Hash-based Message Authentication Code (HMAC) is a hashing algorithm used to verify the integrity and authenticity of messages. Secure Sockets Layer (SSL) requires the use of certificates, so it would require a Public Key Infrastructure (PKI), which is not a simple solution. See Chapter 10.

87. **A.** The best choice is file encryption to protect the passwords in this list. If the passwords were stored in a database, it would be appropriate to encrypt the fields in the database holding the passwords. It's rarely desirable to encrypt an entire database. Whole disk encryption is appropriate for mobile devices. See Chapters 5 and 10.

88. **C.** The most likely issue is that Bart is embedding data in the pictures using steganography techniques. The scenario doesn't give any indications that he is copying the data to a USB drive or encrypting the data, and these actions don't indicate he is leaking the data. If he was sending the data as text in the emails, it would be apparent. See Chapter 10.

89. **A.** Hash-based Message Authentication Code (HMAC) is used with Internet Protocol security (IPsec) and is more likely to be used than any of the

other choices. RFC 4835 mandates the use of HMAC for authentication and integrity. When encryption is used, it also mandates the use of either Advanced Encryption Standard (AES) or Triple Data Encryption Standard (3DES). It does not list Blowfish or Twofish. Message Digest 5 (MD5) is a hashing algorithm. See Chapter 10.

90. **D.** Lisa would decrypt the email with her private key and Bart would encrypt the email with Lisa's public key. Although not part of this scenario, if Bart wanted Lisa to have verification that he sent it, he would create a digital signature with his private key and Lisa would decrypt the private key with Bart's public key. Bart does not need his keys to encrypt email sent to someone else. See Chapter 10.

91. **A.** If Acme submitted the bid via email using a digital signature, it would provide proof that the bid was submitted by Acme. Digital signatures provide verification of who sent a message, non-repudiation preventing them from denying it, and integrity verifying the message wasn't modified. Integrity verifies the message wasn't modified. Repudiation isn't a valid security concept. Encryption protects the confidentiality of data, but it doesn't verify who sent it or provide non-repudiation. See Chapter 10.

92. **C.** Password-Based Key Derivation Function 2 (PBKDF2) is a key stretching technique designed to protect against brute force attempts and is the best choice of the given answers. Another alternative is bcrypt. Both salt the password with additional bits. Triple DES (3DES) is an encryption protocol. Passwords stored using Message Digest 5 (MD5) are easier to crack because they don't use salts. Storing the passwords in encrypted database fields is a possible solution, but just storing them in unencrypted database fields does not protect them at all. See Chapter 10.

93. **A.** A certificate revocation list (CRL) is a list of certificates that a Certificate Authority (CA) has revoked. The CA stores a database repository of revoked certificates and issues the CRL to anyone who requests it. The Online Certificate Status Protocol (OCSP) validates trust with certificates, but only returns short responses such as good, unknown, or revoked. A certificate signing request (CSR) is used to request certificates. See Chapter 10.

94. **B.** Written security policies are management controls. Encryption and the principle of least privilege are technical controls. Change management is an operational control. See Chapter 11.

95. **D.** A separation of duties policy prevents any single person from performing multiple job functions that might allow the person to commit fraud. A mandatory vacation policy is useful to discover fraud committed by an individual, but this scenario clearly indicates this individual controls too many job functions. Although mandatory access control is the strongest access control method available, it doesn't separate job functions. Change management ensures changes are reviewed before being implemented.

96. **B.** A change management policy helps reduce risk associated with making any changes to systems, including updating them. Load balancing can increase the availability associated with an increased load but not with updates. Incident management refers to security incidents. Key management refers to encryption keys.

97. **D.** A successful attack on the social media site resulting in a data breach can expose the passwords and ultimately affect the company application. Users won't use their company credentials to access the social media site, so this doesn't present a risk to the social media site. Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML)-based data format used for SSO on web browsers and it is commonly used on the Internet.

98. **B.** At this stage, the first response is incident identification. The preparation phase is performed before an incident, and includes steps to prevent incidents. After identifying this as a valid incident (malware infection), the next step is escalation and notification and then mitigation steps.

99. **A.** Chain of custody is the primary issue here because the computer was left unattended for several hours. It's difficult to prove that the data collected is the same data that was on the employee's computer when it was confiscated. Data captured from a disk is not volatile so is not an issue in this scenario. The time offset refers to logged times and is not related to this question. Metrics are measurement tools, such as those used to measure the success of a security awareness program.

100. **D.** The best choice of the available answers is to implement a program to increase security awareness, and it could focus on social engineering attacks. A bring your own device (BYOD) policy or an acceptable use policy (AUP) doesn't apply in this scenario. Training is useful, but training users on data handling won't necessarily educate them on social engineering attacks.

Chapter 1

Mastering Security Basics

CompTIA Security+ objectives covered in this chapter:

- 1.3 Explain network design elements and components.**
 - Layered security / Defense in depth
- 2.6 Explain the importance of security related awareness and training.**
 - User habits (Password behaviors)
- 2.9 Given a scenario, select the appropriate control to meet the goals of security.**
 - Confidentiality (Encryption, Access controls, Steganography)
 - Integrity (Hashing, Digital signatures, Certificates, Non-repudiation)
 - Availability (Redundancy, Fault tolerance, Patching)
 - Safety (Fencing, Lighting, Locks, CCTV, Escape plans, Drills, Escape routes, Testing controls)
- 3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques.**
 - Hardening (Password protection)
- 5.1 Compare and contrast the function and purpose of authentication services.**
 - RADIUS, TACACS+, Kerberos, LDAP, XTACACS, SAML, Secure LDAP
- 5.2 Given a scenario, select the appropriate authentication, authorization, or access control.**
 - Identification vs. authentication vs. authorization
 - Authentication (Tokens, Common access card, Smart card, Multifactor authentication, TOTP, HOTP, CHAP, PAP, Single sign-on, Access control)
 - Authentication factors (Something you are, Something you have, Something you know, Somewhere you are, Something you do)
 - Identification (Biometrics, Personal identification verification card, Username)
 - Federation
 - Transitive trust/authentication
- 5.3 Install and configure security controls when performing account management, based on best practices.**
 - Account policy enforcement (Password complexity, Recovery, Lockout, Password history, Password reuse, Password length)
- 6.2 Given a scenario, use appropriate cryptographic methods.**
 - CHAP, PAP

**

Before you dig into some of the details of security, you should have a solid understanding of core security goals. This chapter introduces many of these core goals to provide a big picture of the concepts, and introduces basic risk concepts. Later chapters dig into these topics in more depth. This chapter also covers authentication—how systems and users provide credentials to verify their identity—including authentication used in remote access systems.

Understanding Core Security Goals

Security starts with several principles that organizations include as core security goals. These principles drive many security-related decisions at multiple levels. Understanding these basic concepts helps to give you a solid foundation in security.

Confidentiality, integrity, and availability together form the security triad. Each element is important to address in any security program. Additionally, a well-designed security program addresses other core security principles, such as safety and layered security.

Confidentiality

Confidentiality prevents the unauthorized disclosure of data. In other words, authorized personnel can access the data, but unauthorized personnel cannot access the data. You can ensure confidentiality using several different methods discussed in the following sections.

Encryption

Encryption scrambles data to make it unreadable by unauthorized personnel. Authorized personnel can decrypt the data to access it, but encryption techniques make it extremely difficult for unauthorized personnel to access encrypted data. Chapter 10, “Understanding Cryptography,” covers encryption in much more depth, including commonly used encryption algorithms like Advanced Encryption Standard (AES).

As an example, imagine you need to transmit Personally Identifiable Information (PII), such as medical information or credit card data via email. You wouldn’t want any unauthorized personnel to access this data, but once you click Send, you’re no longer in control of the data. However, if you encrypt the data before you send it, you protect the confidentiality of the data.

Access Controls

Identification, authentication, and authorization combined provide access controls and help ensure that only authorized personnel can access data. Imagine that you want to grant Maggie access to some data, but you don’t want Homer to be able to access the same data. You use access controls to grant and restrict access. The following bullets introduce key elements of access controls:

- **Identification.** Users claim an identity with a unique username. For example, both Maggie and Homer have separate user accounts identified with unique usernames. When Maggie uses her account, she is claiming the identity of her account.
- **Authentication.** Users prove their identity with authentication, such as with a password. For example, Maggie knows her password, but no one else should know it. When she logs on to her account with her username and password, she is claiming the identity of her account and proving her identity with the password. The “Exploring Authentication Concepts” section later in this chapter discusses multiple methods of authentication.
- **Authorization.** Next, you can grant or restrict access to resources using an authorization method, such as permissions. For example, you can grant Maggie’s account full control access to files and folders, and ensure that Homer doesn’t have any permissions to access the data. Chapter 2, “Exploring Control Types and Methods,” digs into access control models a little deeper, including how they enforce authorization settings.

Steganography

A third method you can use for confidentiality is steganography. Chapter 10 covers steganography in more depth, but as an introduction, it is the practice of hiding data within data. Many people refer to it as hiding data in plain sight. For example, you can embed a hidden message in an image by modifying certain bits within the file. If other people look at the file, they won’t notice anything. However, if other people know what to look for, they will be able to retrieve the message.

As a simpler example, you can add a text file to an image file without the use of any special tools other than WinRAR and the Windows command line. If you’re interested in seeing how to do this, check out the Steganography Lab in the online exercises for this book at <http://gcapremium.com/labs/>.

Remember this

Confidentiality ensures that data is only viewable by authorized users. The best way to protect the confidentiality of data is by encrypting it. This includes any type of data, such as PII, data in databases, and data on mobile devices. Access controls help protect confidentiality by restricting access. Steganography helps provide confidentiality by hiding data, such as hiding text files within an image file.

Integrity

Integrity provides assurances that data has not changed. This includes ensuring that no one has modified, tampered with, or corrupted the data. Ideally, only authorized users modify data. However, there are times when unauthorized or unintended changes occur. This can be from unauthorized users, from malicious software (malware), and through system and human errors. When this occurs, the data has lost integrity.

Hashing

You can use hashing techniques to enforce integrity. Chapter 10 discusses the relevant hashing algorithms, such as Message Digest 5 (MD5), Secure Hash Algorithm 1 (SHA-1), and Hash-based Message Authentication Code (HMAC). Briefly, a *hash* is simply a number created by executing a hashing algorithm against data, such as a file or message. As long as the data never changes, the resulting hash will always be the same. By comparing hashes created at two different times, you can determine if the original data is still the same. If the hashes are the same, the data is the same. If the hashes are different, the data has changed.

For example, imagine Homer is sending a message to Marge and they both want assurances that the message retained integrity. Homer's message is, "The price is \$19.99." He creates a hash of this message. For simplicity's sake, imagine the hash is 123. He then sends both the message and the hash to Marge.

Marge receives both the message and the hash. She can calculate the hash on the received message and compare her hash with the hash that Homer sent. If the hash of the received message is 123 (the same as the hash of the sent message), she knows the message hasn't lost data integrity. However, if the hash of the received message is something different, such as 456, then she knows that the received message is not the same. Data integrity has been lost.

Hashing doesn't tell you what modified the message. It only tells you that the message has been modified, with the implication that the information should not be trusted as valid.

You can use hashes with messages, such as email, and any other type of data files. Some email programs use a message authentication code (MAC) instead of a hash to verify integrity, but the underlying concept works the same way.

Acronyms

Don't you just love all of these acronyms? MD5, SHA-1, RAID. There are actually three different meanings of MAC within the context of CompTIA Security+:

1. Media access control (MAC) addresses are the physical addresses assigned to network interface cards (NICs).
2. The mandatory access control (MAC) model is one of several access control models discussed in Chapter 2.
3. Message authentication code (MAC) provides integrity similar to how a hash is used.

If you're having trouble keeping them all straight, don't feel alone. Appendix A, "Acronym List," spells out—and lists brief descriptions—for all of the acronyms used in this book.

...

You can also use hashing techniques to verify that integrity is maintained when files are downloaded or transferred. Some programs can automatically check hashes and determine if a file loses even a single bit during the download process. The program performing the download will detect it by comparing the source hash with the destination hash. If a program detects that the hashes are different, it knows that integrity has been lost and reports the problem to the user.

As another example, a web site administrator can calculate and post the hash of a file on a web site. Users can manually calculate the hash of the file after downloading it and compare the calculated hash with the posted hash. If a virus infects a file on the web server, the hash on the infected file would be different from the hash on the original file (and the hash posted on the web site). You can use freeware such as *md5sum.exe* to calculate MD5 hashes. If you want to see this in action, check out the Hashing Lab in the online exercises for this book at <http://gcgapremium.com/labs/>.

It's also possible to lose data integrity through human error. For example, if a database administrator needs to modify a significant amount of data in a database, the administrator can write a script to perform a bulk update. However, if the script is faulty, it can corrupt the database, resulting in a loss of integrity.

Two key concepts related to integrity are as follows:

- **Integrity provides assurances that data has not been modified, tampered with, or corrupted.** Loss of integrity indicates the data is different. Unauthorized users can change data, or the changes can occur through system or human errors.
- **Hashing verifies integrity.** A hash is simply a numeric value created by executing a hashing algorithm against a message or file. Hashes are created at the source and destination or at two different times (such as on the first and fifteenth of the month). If the hashes are the same, integrity is maintained. If the two hashes are different, data integrity has been lost.

Remember this

Integrity verifies that data has not been modified. Loss of integrity can occur through unauthorized or unintended changes. Hashing algorithms, such as MD5, HMAC, or SHA-1, calculate hashes to verify integrity. A hash is simply a number created by applying the algorithm to a file or message at different times. By comparing the hashes, you can verify integrity has been maintained.

Digital Signatures, Certificates, and Non-Repudiation

You can also use digital signatures for integrity. Chapter 10 covers digital signatures in more depth, but as an introduction, a digital signature is similar in concept to a handwritten signature. Imagine you sign a one-page contract. Anyone can look at the contract later, see your signature, and know it is the same contract. It isn't possible for other people to modify the words in the contract unless they can reproduce your signature, which isn't easy to do.

It's common to use digital signatures with email. For example, imagine that Lisa wants to send an email to Bart. She can attach a digital signature to the email and when Bart receives it, the digital signature provides assurances to him that the email has not been modified.

A digital signature also provides authentication. In other words, if the digital signature arrives intact, it authenticates the sender and Bart knows that Lisa sent it.

Authentication from the digital signature prevents attackers from impersonating others and sending malicious emails. For example, an attacker could make an email look like it came from Lisa and include a link to a malicious web site urging Bart to click it. Without a digital signature, Bart might be fooled into thinking that Lisa sent it and click the link. This might result in Bart inadvertently downloading malware onto his system.

Digital signatures also provide non-repudiation. In other words, Lisa cannot later deny sending the email because the digital signature proves she did. Another way of thinking about non-repudiation is with credit cards. If you buy something with a credit card and sign the receipt, you can't later deny making the purchase. If you do, the store will use your signature to repudiate your claim. In other words, they use your signature for non-repudiation.

Security systems implement non-repudiation methods in other ways beyond digital signatures. Another example is with audit logs that record details such as who, what, when, and where. Imagine Bart logged on to a computer with his username and password, and then deleted several important files. If the audit log recorded these actions, it provides non-repudiation. Bart cannot believably deny he deleted the files.

Digital signatures require the use of certificates and a Public Key Infrastructure (PKI). Certificates include keys used for encryption and the PKI provides the means to create, manage, and distribute certificates. Obviously, there’s much more to certificates and a PKI, but there isn’t room in this chapter for more than an introduction. Feel free to jump ahead to Chapter 10 if you want to learn more right now.

Remember this

Digital signatures can verify the integrity of emails and files. Digital signatures require certificates and also provide authentication and non-repudiation.

Availability

Availability indicates that data and services are available when needed. For some organizations, this simply means that the data and services must be available between 8:00 a.m. and 5:00 p.m., Monday through Friday. For other organizations, this means they must be available 24 hours a day, 7 days a week, 365 days a year.

Organizations commonly implement redundancy and fault-tolerant methods to ensure high levels of availability for key systems. Additionally, organizations ensure systems stay up to date with current patches to ensure that software bugs don't affect their availability.

Redundancy and Fault Tolerance

Redundancy adds duplication to critical systems and provides fault tolerance. If a critical component has a fault, the duplication provided by the redundancy allows the service to continue without interruption. In other words, a system with fault tolerance can suffer a fault, but tolerate it and continue to operate.

A common goal of fault tolerance and redundancy techniques is to remove each single point of failure (SPOF). If an SPOF fails, the entire system can fail. For example, if a server has a single drive, the drive is an SPOF because its failure takes down the server.

Chapter 9, "Preparing for Business Continuity," covers many fault-tolerance and redundancy techniques in more depth. As an introduction, here are some common examples:

- **Disk redundancies.** Fault-tolerant disks such as RAID-1, RAID-5, and RAID-6 allow a system to continue to operate even if a disk fails.
- **Server redundancies.** Failover clusters include redundant servers and ensure a service will continue to operate, even if a server fails. In a failover cluster, the service switches from the failed server in a cluster to an operational server in the same cluster. Virtualization can also increase availability of servers by reducing unplanned downtime. Chapter 5, "Securing Hosts and Data," covers virtualization in more depth.
- **Load balancing.** Load balancing uses multiple servers to support a single service, such as a high-volume web site. It can increase the availability of web sites and web-based applications.
- **Site redundancies.** If a site can no longer function due to a disaster, such as a fire, flood, hurricane, or earthquake, the organization can move critical systems to an alternate site. The alternate site can be a hot site (ready and available 24/7), a cold site (a location where equipment, data, and personnel can be moved to when needed), or a warm site (a compromise between a hot site and cold site).
- **Backups.** If personnel back up important data, they can restore it if the original data is lost. Data can be lost due to corruption, deletion, application errors, human error, and even hungry gremlins that just randomly decide to eat your data. Without data backups, data is lost forever after any one of these incidents.
- **Alternate power.** Uninterruptible power supplies (UPSs) and power generators can provide power to key systems even if commercial power fails.
- **Cooling systems.** Heating, ventilation, and air conditioning (HVAC) systems improve the availability of systems by reducing outages from overheating.

Remember this

Availability ensures that systems are up and operational when needed and often addresses single points of failure. You can increase availability by adding fault tolerance and redundancies, such as RAID, failover clusters, backups, and generators. HVAC systems also increase availability.

Patching

Another method of ensuring systems stay available is with patching. Software bugs cause a wide range of problems, including security issues and even random crashes. When software vendors discover the bugs, they develop and release code that patches or resolves these problems. Organizations commonly implement patch management programs to ensure that systems stay up to date with current patches. Chapter 5 covers patching and patch management in greater depth.

Safety

Another common goal of security is safety. This refers to the safety of both individuals and an organization's assets. You can always replace things, but you cannot replace people, so safety of people should always be a top priority. The following bullets identify some things to consider for both people and assets:

- **Safety of people.** Some of the biggest risks for people occur during disasters, such as fires, earthquakes, hurricanes, and tornadoes. Organizations develop business continuity plans to prepare for these disasters. These plans include items such as escape plans and escape routes. They also ensure personnel are aware of these plans by holding drills and training. You can read more about business continuity in Chapter 9.
- **Safety of assets.** A wide variety of physical security controls helps ensure the safety of assets. These include elements such as fencing around a building, lighting, locks, and closed-circuit television (CCTV) systems to provide video monitoring. Chapter 2 covers physical security controls in more depth.

Organizations test these plans and controls to ensure that they operate as expected. For example, one method of testing controls is to ensure that the escape routes are valid by actually following the route.

Some systems sacrifice security of physical assets to protect people. For example, imagine that the entrance doors to a data center use electronic locks. A fire might result in a power loss, trapping people inside. This isn't acceptable. As an alternative, electronic doors are often designed to fail-open for personnel safety. If the system loses power, the electronic doors will fail in an open state. Unfortunately, if criminals know this, they might decide to destroy the electronic lock or kill power to the building to gain access. Another safety option is to include a method of manually opening the door.

Remember this

Beyond confidentiality, integrity, and availability, safety is another common goal of security. For example, adding fencing and lighting around an organization's property provides safety for personnel and other assets. Similarly, adding stronger locks and door access systems increases safety. Exit doors with electronic locks typically fail in an open position so that personnel can exit safely.

Layered Security/Defense in Depth

Layered security/defense in depth refers to the security practice of implementing several layers of protection. You can't simply take a single action, such as implementing a firewall or installing antivirus software, and consider yourself protected. You must implement security at several different layers. This way, if one layer fails, you still have additional layers to protect you.

If you drive your car to a local Walmart, put a five-dollar bill on the dash, and leave the keys in the car and the car running, there is a very good chance the car won't be there when you come out of the store. On the other hand, if you ensure nothing of value is visible from the windows, the car is locked, and it has an alarm system and stickers on the windows advertising the alarm system, it's less likely that someone will steal it. Not impossible, but less likely.

You've probably heard this as "there is no silver bullet." If you want to kill a werewolf, you can load your gun with a single silver bullet and it will find its mark. The truth is that there is no such thing as a silver bullet. (Of course, there is no such thing as a werewolf either.)

Applied to computers, it's important to implement security at every step, every phase, and every layer. Information technology (IT) professionals can never rest on their laurels with the thought they have done enough and no longer need to worry about security.

It's common to see several layers of protection on the network, such as a firewall, an intrusion detection system (IDS), and proxy servers for content filtering. Chapter 3, "Understanding Basic Network Security," covers these components in more depth, but the key is that they work together to provide several layers of security. Similarly, organizations often install antivirus software on all systems, but also attempt to detect and block malware from entering the network.

Remember this

Layered security, or defense in depth, combines multiple layers of security, such as a firewall, an IDS, content filtering, and antivirus software.

Introducing Basic Risk Concepts

One of the basic goals of implementing IT security is to reduce risk. Because risk is so important and so many chapters refer to elements of risk, it's worth providing a short introduction here.

Risk is the possibility or likelihood of a threat exploiting a vulnerability resulting in a loss. A *threat* is any circumstance or event that has the potential to compromise confidentiality, integrity, or availability. A *vulnerability* is a weakness. It can be a weakness in the hardware, the software, the configuration, or even the users operating the system.

Threats can come from inside an organization, such as from a disgruntled employee, or from outside the organization, such as from an attacker who could be located anywhere on the Internet. Threats can be natural, such as hurricanes, tsunamis, or tornadoes, or man-made, such as malware written by a criminal. Threats can be intentional, such as from attackers, or accidental, such as from employee mistakes or system errors.

Reducing risk is also known as risk mitigation. *Risk mitigation* reduces the chances that a threat will exploit a vulnerability. You reduce risks by implementing controls (also called countermeasures and safeguards), and many of the actions described throughout this book are different types of controls. You can't prevent most threats. For example, you can't stop a tornado or prevent a criminal from writing malware. However, you can reduce risk by reducing vulnerabilities to the threat, or by reducing the impact of the threat.

For example, access controls (starting with authentication) ensure that only authorized personnel have access to specific areas, systems, or data. If employees do become disgruntled and want to cause harm, access controls reduce the amount of potential harm by reducing what they can access. If a natural disaster hits, business continuity and disaster recovery plans help reduce the impact. Similarly, antivirus software prevents the impact of any malware by intercepting it before it causes any harm.

Remember this

Risk is the likelihood that a threat will exploit a vulnerability. Risk mitigation reduces the chances that a threat will exploit a vulnerability, or reduces the impact of the risk, by implementing security controls.

Exploring Authentication Concepts

Authentication proves an identity with some type of credential, such as a username and password. For example, users claim (or profess) their identity with identifiers such as usernames or email addresses. Users then prove their identity with authentication, such as with a password.

In this context, a user’s credentials refer to both a claimed identity and an authentication mechanism. In other words, a user’s credentials can be a username and a password.

At least two entities know the credentials. One entity, such as a user, presents the credentials. The other entity is the authenticator that verifies the credentials. For example, Marge knows her username and password, and an authenticating server knows her username and password. Marge presents her credentials to the authenticating server, and the server authenticates her.

The importance of authentication cannot be understated. You can’t have any type of access control if you can’t identify a user. In other words, if everyone is anonymous, then everyone has the same access to all resources.

Also, authentication is not limited to users. Services, processes, workstations, servers, and network devices all use authentication to prove their identities. Many computers use mutual authentication, where both parties authenticate to each other.

Comparing Identification, Authentication, and Authorization

It's important to understand the differences between identification, authentication, and authorization. When users type in their usernames or email addresses, they are claiming or professing an identity. Users then provide authentication (such as with passwords) to prove their identity.

However, just because users can prove their identity doesn't mean a system automatically grants them access to all the resources. Instead, access control systems grant access to users based on their proven identity. This can be as simple as granting a user permission to read data in a shared folder.

Chapter 2 presents information on different access control models. However, the first step of access control is to implement strong authentication mechanisms, such as the use of complex passwords, smart cards, or biometrics.

Remember this

Identification occurs when a user claims an identity such as with a username or email address.

Authentication occurs when the user proves the claimed identity (such as with a password) and the credentials are verified. Access control systems authorize access to resources based on permissions granted to the proven identity.

Verifying Identities with Identity Proofing

Identity proofing is the process of verifying that people are who they claim to be prior to issuing them credentials, or if they later lose their credentials. Individuals are often required to show other forms of identification, such as their driver's license, before authorities issue the credentials.

This may occur out of view of the IT person creating the account, but identity proofing still occurs. For example, HR personnel process new hires and ensure all their paperwork is in order, including their identification. Later, HR may simply introduce the new employee to an IT professional to create an account. This introduction by the HR person is the only identity proofing needed by the IT worker.

Some identity proofing is electronic. For example, some banking institutions verify your identity based on information about you in credit reports and other available databases. They then ask you a series of multiple-choice questions based on this data, such as your car payment amount, or your height as it's listed on your driver's license.

Identity Proofing for Verification

A second use of identity proofing is after issuing credentials. For example, when a user performs some critical activity (such as transferring money between bank accounts), the bank may reverify the user's identity before transferring the funds.

If you've signed up for online banking, you've probably seen this. When you sign up and sometimes periodically afterward, the bank provides you with a list of questions and records your answers. Some examples are "What is the name of your first pet," "What is the name of your closest childhood friend," "What is the middle name of your oldest sibling," and so on. The idea is that a criminal wouldn't know these answers. If you attempt to transfer money, they might pose these questions to you again before they authorize the transfer.

Many financial institutions record information about the computer you use to access their site. If you later use a different computer, or a computer with a different IP address, the site prompts you to answer one of the identity-proofing questions. If criminals somehow steal your credentials and use them to access your bank account from another computer, identity-proofing techniques help reduce the risk of their success. For example, a criminal from Nigeria shouldn't know these answers. This technique reduces the chances of success for criminals attempting unauthorized transfers, even if they did discover your credentials.

In the past, identity-proofing questions consisted of very few items, such as your birth date, Social Security number, and your mother's maiden name. Because so many entities requested this information and didn't always protect it, it became easy for attackers to obtain this information and use it to steal identities. Some criminals have even stolen identities using information visible on a renter's application.

Unfortunately, many banks have had security breaches compromising their databases (including your answers to these personal questions). In other words, it's becoming less likely that only you know these answers, and they may not be valid as identity-proofing tools for very long.

Self-Service Password Reset Systems

An additional use of identity proofing is with password reset or password recovery systems. These systems provide automated password recovery and are extremely useful in systems with a large number of users. They can actually reduce the total cost of ownership of the system.

Instead of an IT professional spending valuable time resetting passwords, a self-service password reset or password recovery system uses identity proofing to automate the process.

For example, many online systems include a link such as "Forgot Password." If you click on this link, the system might send you your password via email, or reset your password and send the new password via email. Some systems invoke an identity-proofing system. The identity-proofing system asks you questions that you previously provided, such as the name of your first dog, the name of your first boss, and so on. Once you adequately prove your identity, the system gives you the opportunity to change your password.

Of course, an online password reset system won't help a user who can't get online. Some organizations utilize phone-based password reset systems. Users who have forgotten their passwords can call the password reset system and reset their password by using a simple identity-proofing method such as a personal identification number (PIN). A PIN is a simple password, typically four to six numbers long.

Comparing Authentication Factors

Authentication is often simplified as types, or factors, of authentication. Entities can authenticate with any one of these factors, and administrators often combine two factors for dual-factor authentication, and two or more factors for multifactor authentication. As an introduction, the factors are:

- *Something you know*, such as a password or PIN
- *Something you have*, such as a smart card or USB token
- *Something you are*, such as a fingerprint or other biometric identification
- *Somewhere you are*, such as your location using geolocation technologies
- *Something you do*, such as gestures on a touch screen

Something You Know

The *something you know* authentication factor typically refers to a shared secret, such as a password or even a PIN. This factor is the least secure form of authentication. However, you can increase the security of a password by following some simple guidelines. The following sections provide more details, but here's a quick introduction of some of the important password security concepts:

- **Use strong passwords.** This means they are at least eight characters and include multiple character types, such as uppercase letters, lowercase letters, numbers, and symbols.
- **Change passwords regularly.** Systems force users to change their passwords on a regular basis by setting maximum password ages, or password expiration times.
- **Verify a user's identity before resetting a password.** Also, reset the password with a temporary one that expires upon first use.
- **Do not reuse the same passwords.** Password histories prevent users from using the same passwords repeatedly.
- **Implement account lockout policies.** If a user enters the wrong password too many times, an account lockout policy locks the account. This prevents password-guessing attempts.
- **Change default passwords.** If a system comes with a default password, administrators should change it before putting the system into service.
- **Do not write passwords down.** This includes not writing them on Post-it notes and sticking the Post-it notes under a keyboard. If you must write the password down, you should store it in a safe (not just a safe place).
- **Do not share passwords.** Only one person should know the password to any single account. If an administrator resets a password, the password should be set to expire immediately. This requires users to reset the password the first time they log on.

Many systems use technical password policies to enforce many of these guidelines. Chapter 2 covers technical password policies in more depth.

Remember this

The first factor of authentication (*something you know*, such as a password or PIN) is the weakest factor. Passwords should be strong, changed regularly, never shared with another person, and stored in a safe if written down. Technical methods (such as a technical password policy) ensure that users regularly change their passwords and don't reuse the same passwords.

Creating Strong Passwords

One method used to make passwords more secure is to require them to be strong. A strong password is at least eight characters in length, doesn't include words found in a dictionary or any part of a user's name, and combines three of the four following character types:

- Uppercase characters (26 letters A–Z)
- Lowercase characters (26 letters a–z)
- Numbers (10 numbers 0–9)
- Special characters (32 printable characters, such as !, \$, and *)

A complex password uses multiple character types, such as Ab0@. However, a complex password isn't necessarily strong. It also needs to be sufficiently long. It's worth noting that recommendations for the best length of a strong password vary depending on the type of account. A lot of documentation recommends a password length of at least 8 characters for a regular user, and organizations often require administrators to create passwords at least 15 characters long. Additionally, the recommended length is a moving target. Some security experts have been recommending passwords of at least 10 characters for regular users. A key point is that longer passwords are more secure and short passwords of 4 or 5 characters are extremely weak.

The combination of different characters in a password makes up the key space, and you can calculate the key space with the following formula: C^N (CN). C is the number of possible characters used, and N is the length of the password. The ^ character in C^N indicates that C is raised to the N power.

For example, a 6-character password using only lowercase letters (26 letters) is calculated as 26^6 (26^6), or about 308 million possibilities.

Change this to a 10-character password and the value is 26^{10} (26^{10}), or about 141 trillion possibilities. Although this looks like a high number of possibilities, there are password-cracking tools that can test more than 20 billion passwords per second on desktop computers with a high-end graphics processor. An attacker can crack a 10-character password using only lowercase characters (141 trillion possibilities) in less than two hours.

However, if you use all 94 printable characters (uppercase, lowercase, numbers, and special characters) with the same 6- and 10-character password lengths, the values change significantly: 94^6 (94^6) is about 689 billion possibilities, and 94^{10} (94^{10}) is about 53 quintillion. That's 53 followed by 18 zeroes.

You probably don't come across quintillion very often. The order is million, billion, trillion, quadrillion, and then quintillion. The password-cracking tool that cracks a lowercase password in 2 hours will take years to crack a 10-character password using all four character types.

Security experts often mention that if you make a password too complex, you make it less secure.

Read that again. It is not a typo.

More complexity equates to less security. This is because users have problems remembering overly complex passwords such as 4%kiElNsB* and they are more likely to write them down. A password written on paper or stored in a file on a user's computer significantly reduces security.

Instead, users are encouraged to use passphrases. Instead of nonsensical strings of characters, a *passphrase* is a long string of characters that has meaning to the user. A few examples of strong passphrases are ILOveSecurity+, ILOveThi\$B00k, and IW11P@\$\$. Note that these examples include all four character types—uppercase letters, lowercase letters, one or more numbers, and one or more special characters. These passwords are also known as passphrases because they are a combination of words that are easier to remember than a nonsensical string of characters such as 4*eiRS@<].

Strong passwords never include words that can be easily guessed, such as a user's name, words in a dictionary (for any language), or common key combinations.

Remember this

Complex passwords use a mix of character types. Strong passwords use a mix of character types and have a minimum password length of eight characters.

Changing Passwords

In addition to using strong passwords, users should also change their passwords regularly such as every 45 or 90 days. In most systems, technical password policies require users to change their passwords regularly. When the password expires, users are no longer able to log on unless they first change their password.

I can tell you from experience that if users are not forced to change their passwords through technical means, they often simply don't. It doesn't matter how many reminders you give them. On the other hand, when a password policy locks out user accounts until they change their password, they will change it right away.

Resetting Passwords

It's not uncommon for users to occasionally forget their password. In many organizations, help-desk professionals or other administrators reset user passwords.

Before resetting the password, it's important to verify the user's identity. Imagine that Hacker Harry calls into the help desk claiming to be the CEO and asks for *his* password to be reset. If the help-desk professional does so, it locks the CEO out of the account. Worse, depending on the process, it might give Hacker Harry access to the CEO's account. Organizations use a variety of different methods of identification, including methods discussed in the "Identity Proofing for Verification" section earlier in this chapter.

In some systems, help-desk professionals manually change the user's password. This causes a different problem. Imagine a user calls the help desk and asks for a password reset. The help-desk professional changes the password and lets the user know the new password. However, at this point two people know the password. The help-desk professional could use the password and impersonate the user, or the user could blame the help-desk professional for impersonating the user.

Instead, the help-desk professional should set the password as a temporary password that expires upon first use. This requires the user to change the password immediately after logging on and it maintains password integrity.

Remember this

Before resetting passwords for users, it's important to verify the user's identity. When resetting passwords manually, it's best to create a temporary password that expires upon first use.

Using Password History

Many users would prefer to use the same password forever simply because it's easier to remember. Even when technical password policies force users to change their passwords, many users simply change them back to the original password. Unfortunately, this significantly weakens password

security.

A password history system remembers past passwords and prevents users from reusing passwords. It's common for password policy settings to remember the last 24 passwords and prevent users from reusing these until they've used 24 new passwords.

When implementing password history, it's best to include a minimum password age setting. For example, a minimum password age of 1 prevents users from changing their password until one day has passed. This prevents users from changing their passwords multiple times to get back to their original password. Chapter 2 shows how to implement password history with a minimum password age.

Remember this

You can combine password history with a minimum password age to prevent users from reusing the same passwords. A password history of 24 remembers the last 24 passwords.

Implementing Account Lockout Policies

Accounts will typically have lockout policies preventing users from guessing the password. If a user enters the wrong password too many times (such as three or five times), the system locks the user's account. Two key phrases associated with account lockout policies are:

- **Account lockout threshold.** This is the maximum number of times a user can enter the wrong password. When the user exceeds the threshold, the system locks the account.
- **Account lockout duration.** This indicates how long an account remains locked. It could be set to 30, indicating that the system will lock the account for 30 minutes. After 30 minutes, the system automatically unlocks the account. If the duration is set to 0, the account remains locked until an administrator unlocks it.

Changing Default Passwords

In Chapter 5, you'll learn the basics of hardening systems, including changing defaults, removing unnecessary protocols and services, and keeping the system up to date.

Many systems and devices have default passwords. A basic security practice is to change these defaults before putting a system into use. As an example, many wireless routers have default accounts named "admin" with a default password of "admin." If you don't change these defaults, anyone who knows the defaults can log on and take control of the router. In that case, the attacker can even go as far as locking you out of your own network.

Changing defaults also includes changing the default name of the Administrator account, if possible. In many systems, the Administrator account can't be locked out through regular lockout policies, so an attacker can continue to try to guess the password of the Administrator account without risking being locked out. By changing the name of the Administrator account to something else such as Not4U2Know, it reduces the chances of success for the attacker. The attacker needs to know the new administrator name before he can try to guess the password.

Some administrators go a step further and add a dummy user account named "administrator." This account has no permissions. If someone does try to guess the password of this account, the system will lock it out, alerting administrators of possible illicit activity.

Using Previous Logon Notification

A simple technique used to alert users of possible account problems is to provide them notification of when they last logged on. You might see this as "Previous logon notification" when you first log on to a system.

As an example, consider Maggie, who took Friday off last week. She logged on last Thursday while she was at work, but she didn't log on again all weekend. When she came in to work on Monday, she logged on and the system notified her that the last time she logged on was on Friday. If she's paying attention to this message, she'll realize that someone else logged on to her account. This also alerts her that her credentials have been compromised.

The primary challenge with this system is that users tend to ignore the notification. More than 99 percent of the time, the message tells users what they already know. So, instead of reading the message, users tend to ignore it.

Storing Passwords

You should not write down passwords unless absolutely necessary. If you do write down passwords, you should store them in a safe. Note that this is not simply a safe place.

Many users have simply written down their passwords on a Post-it note and stuck it to the bottom of their keyboard, thinking no one would ever look there. Hackers, crackers, attackers, and even curious fellow employees who have physical access to your system will think to look under a keyboard. At one organization, this was so prevalent that when a computer was disposed of, the computer sanitization checklist included checking under the keyboard.

Sharing Passwords

Only one person should know the password, and users should not share their passwords with anyone. This can be a difficult message to ingrain in the minds of end users, resulting in many successful social engineering attempts.

Chapter 6, “Understanding Malware and Social Engineering,” covers social engineering in more depth, but, for now, be aware that attackers often gain information just by asking. They can ask over the phone, in person, or via email with increasingly sophisticated phishing attacks.

Social engineers use trickery and conniving to convince users to give out their passwords. If administrators train users that it is sometimes OK to share a password, such as when an administrator asks for it, it sets the user up for trouble. A social engineer will find it easier to trick the user into giving up a password. On the other hand, if users consistently hear the message that they should *never* share passwords, alarm bells will ring in a user’s head when a social engineer starts asking for a password.

Training Users About Password Behaviors

Common user habits related to password behaviors have historically ignored security. Many users don’t understand the value of their password, or the potential damage if they give it out. It’s important for an organization to provide adequate training to users on password security if they use passwords within the organization. This includes both the creation of strong passwords and the importance of never giving out their passwords.

For example, the password “123456” frequently appears on lists as the most common password in use. The users who are creating this password probably don’t know that it’s almost like using no password at all. Also, they probably don’t realize that they can significantly increase the password strength by using a simple passphrase such as “IC@nC0untTo6.” A little training can go a long way.

Something You Have

The *something you have* authentication factor refers to something you can physically hold. This section covers many of the common items in this factor, including smart cards, Common Access Cards, and hardware tokens. It also covers two newer open source protocols used with hardware tokens.

Smart Cards

Smart cards are credit card-sized cards that have an embedded microchip and a certificate. Users insert the smart card into a smart card reader, similar to how someone would insert a credit card into a credit card reader. The smart card reader reads the information on the card, including the details from the certificate.

The embedded certificate allows the use of a complex encryption key and provides much more secure authentication than is possible with a simple password. Additionally, the certificate can be used with digital signatures and data encryption. The smart card provides confidentiality, integrity, authentication, and non-repudiation.

Requirements for a smart card are:

- **Embedded certificate.** The embedded certificate holds a user’s private key (which is only accessible to the user) and is matched with a public key (that is publicly available to others). The private key is used each time the user logs on to a network.
- **Public Key Infrastructure (PKI).** Chapter 10 covers PKI in more depth, but in short, the PKI supports issuing and managing certificates.

Smart cards are often used with another factor of authentication. For example, a user may also enter a PIN or password, in addition to using the smart card. Because the smart card is in the *something you have* factor and the PIN is in the *something you know* factor, this combination is dual-factor authentication.

CACs and PIVs

A *Common Access Card (CAC)* is a specialized type of smart card used by the U.S. Department of Defense. In addition to including the capabilities of a smart card, it also includes a picture of a user and other readable information. Users can use the CAC as a form of photo identification to gain access into a secure location. For example, they can show their CAC to guards who are protecting access to secure areas. Once inside the secure area, users can use the CAC as a smart card to log on to computers.

Similarly, a *Personal Identity Verification (PIV)* card is a specialized type of smart card used by U.S. federal agencies. It also includes photo identification and provides confidentiality, integrity, authentication, and non-repudiation for the users, just as a CAC does.

CACs and PIVs both support dual-factor authentication (sometimes called two-factor authentication) because users generally log on with the smart card and by entering information they know such as a password. Additionally, these cards include embedded certificates used for digital signatures and encryption.

Remember this

Smart cards are often used with dual-factor authentication where users have something (the smart card) and know something (such as a password or PIN). Smart cards include embedded certificates used with digital signatures and encryption. CACs and PIVs are specialized smart cards that include photo identification. They are used to gain access into secure locations and to log on to computer systems.

Tokens or Key Fobs

A *token* or *key fob* (sometimes simply called a fob) is an electronic device about the size of a remote key for a car. You can easily carry them in a

pocket or purse, or connect them to a key chain. They include liquid crystal display (LCD) that displays a number that changes periodically, such as every 60 seconds. To differentiate them from logical tokens, they are sometimes called hardware tokens.

The token is synced with a server that knows what the number is at any moment. For example, at 9:01, the number displayed on the token may be 135792 and the server knows the number is 135792. At 9:02, the displayed number changes to something else and the server also knows the new number.

This number is a one-time use, rolling password. It isn't useful to attackers for very long, even if they can discover it. For example, a shoulder surfing attacker may be able to look over someone's shoulder and read the number. However, the number expires within the next 60 seconds and is replaced by another one-time password.

Users often use tokens to authenticate via a web site. They enter the number displayed in the token along with their username and password. This provides dual-factor authentication because the user must have something (the token) and know something (their password).

RSA sells RSA Secure ID, a popular token used for authentication. You can Google "Secure ID picture" to view many pictures of these tokens. Although RSA tokens are popular, other brands are available.

USB tokens include a USB connector and a smart chip. The smart chip typically stores a certificate similar to how smart cards store a certificate. In addition to being used for authentication, the embedded certificate supports the use of digital signatures.

HOTP and TOTP

[HMAC-based One-Time Password](#) (HOTP) is an open standard used for creating one-time passwords, similar to those used in tokens or key fobs. The algorithm combines a secret key and an incrementing counter, and then uses HMAC to create a hash of the result. It then converts the result into an HOTP value of six to eight digits.

Imagine Bart needs to use HOTP for authentication. He requests a new HOTP number using a token or a software application. He can then use this number for authentication along with some other authentication method, such as a username and password. As soon as he uses it, the number expires. No one else is able to use it, and Bart cannot use it again either.

Here's an interesting twist, though. A password created with HOTP remains valid until it's used. Suppose Bart requested the HOTP number but then got distracted and never used it. What happens now? Theoretically, it remains usable forever. This presents a risk related to HOTP because other people can use the password if they discover it.

A Time-based One-Time Password (TOTP) is similar to HOTP, but it uses a timestamp instead of a counter. Moreover, one-time passwords created with TOTP expire after 30 seconds.

One significant benefit of HOTP and TOTP is price. Hardware tokens that use these open source standards are significantly less expensive than tokens that use proprietary algorithms. Additionally, many software applications use these algorithms and they are freely available.

For example, Figure 1.1 shows the free VIP Access app created by Symantec and running on an iPad. It's also available for many other tablets and smartphones. Once you configure it to work with a compatible authentication server, it creates a steady stream of one-time use passwords. The six-digit security code is the password, and the counter lets you know how much more time you have before it changes again.




Figure 1.1: VIP Access app

Similar to a hardware token, the user enters a username and password as the *something you know* factor, and then enters the security code from the app as the *something you have* factor. This provides dual-factor authentication. Many public web sites like eBay and PayPal support it, allowing many end users to implement dual-factor authentication as long as they have a smartphone or tablet device.

Remember this

HOTP and TOTP are both open source standards used to create one-time use passwords. HOTP creates a one-time use password that does not expire. TOTP creates a one-time password that expires after 30

seconds.

Something You Are

The *something you are* authentication factor uses biometrics for authentication. Biometric methods are the strongest form of authentication because they are the most difficult for an attacker to falsify. In comparison, passwords are the weakest form of authentication.

Biometric Methods

Biometrics use a physical characteristic, such as a fingerprint, for authentication. Biometric systems use a two-step process. In the first step, users register with the authentication system. For example, an authentication system first captures a user's fingerprint and associates it with the user's identity. Later, when users want to access the system, they use their fingerprints to prove their identity. There are multiple types of biometrics, including:

- **Fingerprint and thumbprint.** Law enforcement agencies have used these for decades. Many laptop computers include fingerprint scanners or fingerprint readers, and they have also begun to appear on tablet devices and smartphones. Similarly, some USB flash drives include a fingerprint scanner. They can store multiple fingerprints of three or four people to share access to the same USB drive.
- **Handprint.** These are similar to fingerprints, but the scanners look at the whole hand. Many amusement parks sell annual passes, but they don't want these passes shared with everyone in the neighborhood. They use biometric hand scanners to authenticate individuals as the actual owner of a pass. If someone else tries to use the pass, the scan fails.
- **Palm.** Palm scanners identify individuals using infrared scanners and a palm-vein pattern recognition system. A great benefit of palm scanners is that you do not need to touch the scanner, eliminating a concern with many other types of scanners. Many businesses outside of the United States are using palm scanners in businesses and fast-food restaurants. Within the United States, many hospitals have begun to use them. Vendors claim that palm-scanning systems are as much as 100 times more accurate than fingerprint-scanning systems.
- **Retina.** Retina scanners scan the retina of one or both eyes and use the pattern of blood vessels at the back of the eye for recognition. Some people object to the use of these scanners for authentication because they can identify medical issues, and because you typically need to have physical contact with the scanner.
- **Iris.** Iris scanners use camera technologies to capture the patterns of the iris around the pupil for recognition. They are used in many passport-free border crossings around the world. They can take pictures from about 3 to 10 inches away, avoiding physical contact.

Although the use of DNA is possible in the future for authentication, it's unlikely it'll be used in the near term. Besides the lack of ability to identify DNA in a timely manner, most users will likely balk at having to prick their fingers to provide a blood sample to authenticate to a computer.

Remember this

The third factor of authentication (*something you are*, defined with biometrics) is the strongest individual method of authentication because it is the most difficult for an attacker to falsify. Biometric methods include fingerprints, retina scans, and palm scanners.

Biometric Errors

Biometrics can be very exact when the technology is implemented accurately. However, it is possible for a biometric manufacturer to take shortcuts and not implement it correctly, resulting in false readings. Two biometric false readings are:

- **False acceptance.** This is when a biometric system incorrectly identifies an unauthorized user as an authorized user. The False Accept Rate (FAR, also known as a type 2 error) identifies the percentage of times false acceptance occurs.
- **False rejection.** This is when a biometric system incorrectly rejects an authorized user. The False Reject Rate (FRR, also known as a type 1 error) identifies the percentage of times false rejections occur.

True readings occur when the biometric system accurately accepts or rejects a user. For example, *true acceptance* is when the biometric system accurately determines a positive match. In contrast, *true rejection* occurs when the biometric system accurately determines a nonmatch.

As an example, a local grocery store had problems with false acceptance. It allowed shoppers to register their debit cards with their fingerprints. Once the shoppers registered their debit cards, they could simply place one of their fingers in a fingerprint reader instead of swiping a debit card and entering a PIN. Then one day, all the fingerprint scanners disappeared. Store management had discovered that the scanners were falsely accepting some users. Users who had not registered their debit cards tried the fingerprint scanners to pay for groceries, and it worked. Instead of rejecting these unknown users, the system falsely accepted them as a known user and charged their grocery purchase to someone else's bank account.

Amusement parks often have problems with false rejection later in the season. Families are often the biggest customers of these annual passes. Children grow enough during the season that the handprint recorded for them early in the season no longer matches for them later in the season.

Somewhere You Are

The *somewhere you are* authentication factor identifies a user's location. *Geolocation* is a group of technologies used to identify a user's location

and is the most common method used in this factor. Many authentication systems use the Internet Protocol (IP) address for geolocation. The IP address provides information on the country, region, state, city, and sometimes even the zip code.

As an example, I once hired a virtual assistant in India to do some data entry for me. I created an account for the assistant in an online application called Hootsuite and sent him the logon information. However, when he attempted to log on, Hootsuite recognized that his IP was in India but I always logged on from an IP in the United States. Hootsuite blocked his access and then sent me an email saying that someone from India was trying to log on. They also provided me directions on how to grant him access if he was a legitimate user, but it was comforting to know they detected and blocked this access automatically.

Within an organization, it's possible to use the computer name or the MAC address of a system for the *somewhere you are* factor. For example, in a Microsoft Active Directory domain, you can configure accounts so that users can only log on to the network through one specific computer. If they aren't at that computer, the system blocks them from logging on at all.

Something You Do

The *something you do* authentication factor refers to actions you can take such as gestures on a touch screen. As an example, Microsoft Windows 8 supports picture passwords. Users first select a picture, and then they can add three gestures as their picture password. Gestures include tapping in specific places on the picture, drawing lines between items with a finger, or drawing a circle around an item such as someone's head. After registering the picture and their gestures, users repeat these gestures to log on again later.

Other examples of *something you do* include how you write or how you type. For example, keystroke dynamics measures the pattern and rhythm as a user types on a keyboard. It measures details such as speed, dwell time, and flight time. *Dwell time* is the time a key is pressed, and *flight time* is the time between releasing one key and pressing the next key. Many security professionals refer to this as behavioral biometrics because it identifies behavioral traits of an individual. However, some people put these actions into the *something you do* authentication factor.

Dual-Factor and Multifactor Authentication

Dual-factor authentication (sometimes called two-factor authentication) uses two different factors of authentication such as *something you have* and *something you know*. Dual-factor authentication often uses a smart card and a PIN, a USB token and a PIN, or combining a smart card or hardware token with a password. In each of these cases, the user must have something and know something.

Multifactor authentication uses two or more factors of authentication. For example, you can combine the *something you are* factor with one or more other factors of authentication.

Note that technically you can call an authentication system using two different factors either dual-factor authentication or multifactor authentication. Multifactor authentication indicates multiple factors and multiple is simply more than one.

It's worth noting that using two methods of authentication in the same factor is not dual-factor authentication. For example, requiring users to enter a password and a PIN (both in the *something you know* factor) is single-factor authentication, not dual-factor authentication. Similarly, using a thumbprint and a retina scan is not dual-factor authentication.

Remember this

Two or more methods in the same factor of authentication (such as a PIN and a password) is single-factor authentication. Dual-factor (or two-factor) authentication uses two *different* factors such as a USB token and a PIN. Multifactor authentication uses two or more factors.

Summarizing Identification Methods

This chapter presented several different identification methods and because identification is so important, it's worthwhile summarizing them. They are usernames, photo identification cards, and biometrics.

The most commonly used identification method is a username. This can be a traditional username such as DarrilGibson and it can be an email address such as Darril@gcgapremium.com depending on how the system is configured. Many other identification methods can be used for both identification and authentication.

CACs and PIVs include a picture and other information about the owner, so owners often use them for identification. They also function as smart cards in the *something you have* authentication factor.

The "Something You Are" section focused on using biometrics for authentication, but several entities also use biometric methods for identification. For example, law enforcement agencies have used fingerprints to identify individuals at crime scenes for decades. Similarly, retina and palm scanners can identify individuals with a high degree of accuracy.

Comparing Authentication Services

Several other authentication services are available that fall outside the scope of the three previously described factors of authentication. The following sections describe many of these services.

Kerberos

Kerberos is a network authentication mechanism used within Windows Active Directory domains and some Unix environments known as realms. It was originally developed at MIT (the Massachusetts Institute of Technology) for Unix systems and later released as a request for comments (RFC). Kerberos provides mutual authentication that can help prevent man-in-the-middle attacks and uses tickets to help prevent replay attacks.

Kerberos includes several requirements for it to work properly. They are:

- **A method of issuing tickets used for authentication.** The Key Distribution Center (KDC) uses a complex process of issuing ticket-granting tickets (TGTs) and other tickets. The KDC (or TGT server) packages user credentials within a ticket. Tickets provide authentication for users when they access resources such as files on a file server. These tickets are sometimes referred to as tokens, but they are logical tokens, not a key-fob type of token discussed in the “Something You Have” section.
- **Time synchronization.** Kerberos version 5 requires all systems to be synchronized and within five minutes of each other. The clock that provides the time synchronization is used to timestamp tickets, ensuring they expire correctly. This helps prevent replay attacks. In a replay attack, a third party attempts to impersonate a client after intercepting data captured in a session. However, if an attacker intercepts a ticket, the timestamp limits the amount of time an attacker can use the ticket.
- **A database of subjects or users.** In a Microsoft environment, this is Active Directory, but it could be any database of users.

When a user logs on with Kerberos, the KDC issues the user a ticket-granting ticket, which typically has a lifetime of 10 hours to be useful for a single workday. When the user tries to access a resource, the ticket-granting ticket is presented as authentication, and the user is issued a ticket for the resource. However, the ticket expires if users stay logged on for an extended period, such as longer than 10 hours. This prevents them from accessing network resources. In this case, users may be prompted to provide a password to renew the ticket-granting ticket, or they may need to log off and back on to generate a new ticket-granting ticket.

Remember this

Kerberos is a network authentication protocol within a Microsoft Windows Active Directory domain or a Unix realm. It uses a database of objects such as Active Directory and a KDC (or TGT server) to issue timestamped tickets that expire after a certain time period.

Additionally, Kerberos uses symmetric-key cryptography to prevent unauthorized disclosure and to ensure confidentiality. Chapter 10 explains algorithms in more depth, but in short, symmetric-key cryptography uses a single key for both encryption and decryption of the same data.

In contrast, asymmetric encryption uses two keys: one key to encrypt and one key to decrypt. Asymmetric encryption requires a PKI to issue certificates. The two keys used in a PKI are a public key and a private key created as matched pairs. Information encrypted with the public key can only be decrypted with the matching private key. Similarly, information encrypted with the private key can only be decrypted with the matching public key.

As a memory trick, you might want to remember this: Symmetric encryption uses one key. Asymmetric adds a syllable (“a”), and it also adds a key, using two keys.

LDAP and Secure LDAP

Lightweight Directory Access Protocol (LDAP) specifies formats and methods to query directories. In this context, a *directory* is a database of objects that provides a central access point to manage users, computers, and other directory objects. LDAP is an extension of the X.500 standard that Novell and early Microsoft Exchange Server versions used extensively.

Windows domains use Active Directory, which is based on LDAP. *Active Directory* is a directory of objects (such as users, computers, and groups), and it provides a single location for object management. Queries to Active Directory use the LDAP format. Similarly, Unix realms use LDAP to identify objects.

Administrators often use LDAP in scripts, but they need to have a basic understanding of how to identify objects. For example, a user named Homer in the Users container within the GetCertifiedGetAhead.com domain is identified with the following LDAP string:

LDAP://CN=Homer,CN=Users,DC=GetCertifiedGetAhead,DC=com

- **CN=Homer.** CN is short for common name.
- **CN=Users.** CN is sometimes referred to as container in this context.
- **DC=GetCertifiedGetAhead.** DC is short for domain component.
- **DC=com.** This is the second domain component in the domain name.

Secure LDAP uses encryption to protect LDAP transmissions. When a client connects with a server using Secure LDAP, the two systems establish a Transport Layer Security (TLS) session before transmitting any data. TLS encrypts the data before transmission.

LDAP Version 2 (LDAP v2) uses Secure Sockets Layer (SSL) instead of TLS. However, LDAP Version 3 (LDAP v3) is the current standard and it uses TLS.

Remember this

LDAP is based on an earlier version of X.500. Windows Active Directory domains and Unix realms use LDAP to identify objects in query strings with codes such as CN=Users and DC=GetCertifiedGetAhead. Secure LDAP encrypts transmissions with SSL or TLS.

Single Sign-On

Single sign-on (SSO) refers to the ability of a user to log on or access multiple systems by providing credentials only once. SSO increases security because the user only needs to remember one set of credentials and is less likely to write them down. It's also much more convenient for users to access network resources if they only have to log on one time.

As an example, consider a user who needs to access multiple servers within a network to perform normal work. Without SSO, the user would need to know one set of credentials to log on locally, and additional credentials for each of the servers. Many users would write these credentials down to remember them.

Alternatively, in a network with SSO capabilities, the user only needs to log on to the network once. The SSO system typically creates some type of SSO token used during the entire logon session. Each time the user accesses a network resource, the SSO system uses this token for authentication. Kerberos and LDAP both include SSO capabilities.

Remember this

Single sign-on enhances security by requiring users to use and remember only one set of credentials for authentication. Once signed on using SSO, this one set of credentials is used throughout a user's entire session. SSO can provide central authentication against a federated database for different operating systems.

Same sign-on is not the same as SSO. In a same sign-on system, users have to reenter their credentials each time they access another system. However, they use the same credentials.

SSO and Transitive Trusts

A transitive trust creates an indirect trust relationship. As an example, imagine a transitive trust relationship exists between Homer, Moe, and Fat Tony:

- Homer trusts Moe.
- Moe trusts Fat Tony.
- Because of the transitive trust relationship, Homer trusts Fat Tony.

Of course, this isn't always true with people and Homer may be a little upset with Moe if he shares his secrets with Fat Tony. However, it reduces network administration in a domain.

Within an LDAP-based network, domains use transitive trusts for SSO. Figure 1.2 shows a common configuration with three domains in the same network.

The parent domain is GetCertifiedGetAhead.com and the configuration includes two child domains—Training and Blogs.

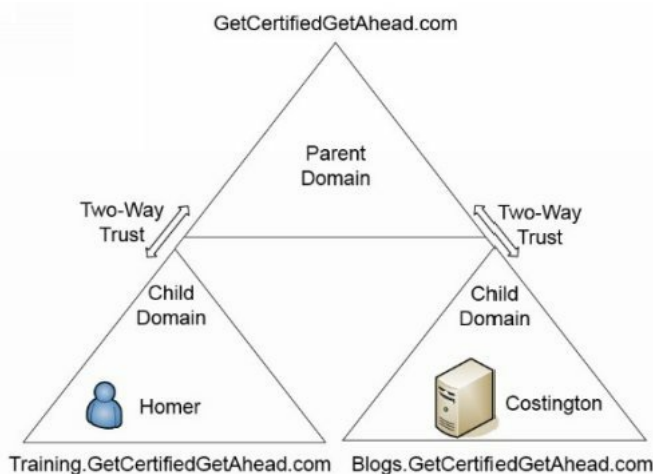


Figure 1.2: An LDAP transitive trust used for SSO

In this example, there is a two-way trust between the parent domain and the child domain, GetCertifiedGetAhead.com and Training.GetCertifiedGetAhead.com, respectively. The parent trusts the child, and the child trusts the parent. Similarly, there is a two-way trust between the parent domain and the Blogs child domain. There isn't a direct trust between the two child domains. However, the transitive relationship creates a two-way trust between them.

All of these domains contain objects such as users, computers, and groups. Homer's user account is in the Training domain, and a server named Costington is in the Blogs domain. With the transitive trust, it's possible to grant Homer access to the Costington server without creating another trust relationship directly between Training and Blogs.

Without a trust relationship, you'd have to create another account for Homer in the Blogs domain before you could grant him access. Additionally, Homer would need to manage the second account's password separately. However, with the transitive trust relationships, the network supports SSO so

Homer only needs a single account.

SSO and a Federation

Some SSO systems can connect authentication mechanisms from different environments, such as different operating systems or different networks. One common method is with a federated identity management system, often integrated as a federated database. This federated database provides central authentication in a nonhomogeneous environment.

As an example, imagine that the Springfield Nuclear Power Plant established a relationship with the Springfield school system, allowing the power plant employees to access school resources. It's not feasible or desirable to join these two networks into one. However, you can create a federation of the two networks. Once it's established, the power plant employees will log on using their power plant account, and then access the shared school resources without logging on again.

A federation requires a federated identity management system that all members of the federation use. In the previous example, the members of the federation are the power plant and the school system. Members of the federation agree on a standard for federated identities and then exchange the information based on the standard. A federated identity links a user's credentials from different networks or operating systems, but the federation treats it as one identity.

SSO and SAML

Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML)–based data format used for SSO on web browsers. Imagine two web sites hosted by two different organizations. Normally a user would have to provide different credentials to access either web site. However, if the organizations trust each other, they can use SAML as a federated identity management system. Users authenticate with one web site and are not required to authenticate again when accessing the second web site.

Many web-based portals use SAML for SSO. The user logs on to the portal once, and the portal then passes proof of the user's authentication to back-end systems. As long as one organization has authenticated users, they are not required to authenticate again to access other sites within the portal.

SAML defines three roles:

- **Principal.** This is typically a user. The user logs on once. If necessary, the principal requests an identity from the identity provider.
- **Identity provider.** An identity provider creates, maintains, and manages identity information for principals.
- **Service provider.** A service provider is an entity that provides services to principals. For example, a service provider could host one or more web sites accessible through a web-based portal. When a principal tries to access a resource, the service provider redirects the principal to obtain an identity first.

This process sends several XML-based messages between the systems. However, it is transparent to the user.

SAML and Authorization

It's important to realize that the primary purpose of SSO is for identification and authentication of users. Users claim an identity and prove that identity with credentials. SSO does not provide authorization. For example, if the power plant and the school system create a federation using SAML, this doesn't automatically grant everyone in the school system full access to the nuclear power plant resources. Authorization is completely separate.

However, many federation SSO systems, including SAML, include the ability to transfer authorization data between their systems. In other words, it's possible to use SAML for single sign-on authentication and for authorization.

Remember this

SAML is an XML-based standard used to exchange authentication and authorization information between different parties. SAML provides SSO for web-based applications.

Authenticating RAS Clients

Remote Access Service (RAS) provides access to an internal network from an outside source. Chapter 4, “Securing Your Network,” covers RAS in more depth, but this section covers different authentication mechanisms you can use with RAS.

Clients access a RAS server via either dial-up or a virtual private network (VPN). A VPN allows a client to access a private network over a public network, such as the Internet.

Remote access methods are useful for personnel who need access to the private network from remote locations. This includes users who travel frequently and telecommuters who work from home. However, no matter which remote access method you use, you still need to ensure that only authorized clients can access the network remotely.

Authorization begins with authentication, and there are multiple methods of authentication used with remote access. The following sections describe the different remote access authentication mechanisms in more depth, but here’s a quick introduction:

- **Password Authentication Protocol (PAP).** PAP sends passwords in cleartext so PAP is used only as a last resort.
- **Challenge Handshake Authentication Protocol (CHAP).** CHAP uses a handshake process where the server challenges the client. The client then responds with appropriate authentication information.
- **Microsoft CHAP (MS-CHAP).** This is the Microsoft implementation of CHAP, which is used only by Microsoft clients.
- **MS-CHAPv2.** MS-CHAP is deprecated in favor of MS-CHAPv2. It includes several improvements, including the ability to perform mutual authentication.
- **Remote Authentication Dial-In User Service (RADIUS).** RADIUS provides a centralized method of authentication for multiple remote access servers. RADIUS encrypts the password packets, but not the entire authentication process.
- **Diameter.** Diameter is an improvement over RADIUS and it supports Extensible Authentication Protocol (EAP) for security.
- **Extended Terminal Access Controller Access-Control System (XTACACS).** Cisco Systems developed XTACACS as an improvement over TACACS and it is proprietary to Cisco systems.
- **Terminal Access Controller Access-Control System Plus (TACACS+).** TACACS+ is an alternative to RADIUS, but it is proprietary to Cisco systems. A benefit of TACACS+ is that it can interact with Kerberos, allowing it to work with a broader range of environments, including Microsoft domains using Kerberos. Additionally, TACACS+ encrypts the entire authentication process, whereas RADIUS encrypts only the password.

PAP

Password Authentication Protocol (PAP) is used with Point-to-Point Protocol (PPP) to authenticate clients. It replaced Serial Line Interface Protocol (SLIP) as a more efficient method of connecting to remote servers such as Internet Service Providers (ISPs). However, a significant weakness of PAP is that it sends passwords or PINs over a network in cleartext, representing a significant security risk.

PPP was primarily used with dial-up connections. Believe it or not, there was a time when the thought of someone wiretapping a phone was rather remote. Because of this, security was an afterthought with PPP. Today, PPP is only used as a last resort due to passwords being passed in cleartext, or it is used with another protocol that provides encryption.

Throughout this book, you'll read that sending data across a network in cleartext is a security risk. This is a risk with PPP and with many other protocols such as File Transfer Protocol (FTP). It's relatively easy for attackers to download and use a free protocol analyzer (such as Wireshark) to capture packets. They can then analyze these packets, and read the data within the packets. Protocol analyzers are commonly referred to as sniffers, and Chapter 8, "Managing Risk," covers them in more depth.

CHAP

Challenge Handshake Authentication Protocol (CHAP) also uses PPP and authenticates remote users, but it is more secure than PAP. The goal of CHAP is to allow the client to pass credentials over a public network (such as a phone or the Internet) without allowing attackers to intercept the data and later use it in an attack.

The client and server both know a shared secret (similar to a password) used in the authentication process. However, the client doesn't send the shared secret over the network in plaintext as PAP does. Instead, the client hashes it after combining it with a nonce (number used once) provided by the server. This handshake process is used when the client initially tries to connect to the server, and at different times during the connection.

Remember this

PAP authentication uses a password or a PIN. A significant weakness is that PAP sends the information across a network in cleartext, making it susceptible to sniffing attacks. CHAP is more secure than PAP because passwords are not sent over the network in cleartext. Both PAP and CHAP use PPP.

MS-CHAP and MS-CHAPv2

Microsoft introduced Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) as an improvement over CHAP for Microsoft clients. MS-CHAP supported clients as old as Windows 95. Later, Microsoft improved MS-CHAP with MS-CHAPv2.

A significant improvement of MS-CHAPv2 over MS-CHAP is the ability to perform mutual authentication. Not only does the client authenticate to the server, but the server also authenticates to the client. Chapter 7, “Identifying Advanced Attacks,” covers different types of attacks, including attacks in which an attacker may try to impersonate a server. Mutual authentication provides assurances of the server’s identity before the client transmits data, which reduces the risk of a client sending sensitive data to a rogue server.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a centralized authentication service. Instead of each individual RAS server needing a separate database to identify who can authenticate, authentication requests are forwarded to a central RADIUS server.

One of the ways to visualize this is to think of a dial-up ISP, such as AOL. AOL provides dial-up digital subscriber line (DSL) services for many cities in the United States. Although anyone can dial in to an AOL server, the servers only allow access to users with accounts in good standing.

Suppose you live in Virginia Beach, Virginia. You could sign up for AOL, pay the required fee, and access the AOL service via a server in Virginia Beach. AOL could maintain a database on the server in Virginia Beach showing your account is in good standing.

However, what if you travel to Atlanta, Chicago, San Francisco, or somewhere else? Because you paid your fees, you would reasonably expect to be able to access AOL no matter where you traveled. If the databases were stored on individual servers, AOL would need to update every single server in the United States. Similarly, if you canceled your subscription, AOL would need to update every single server again. Clearly, this would take a lot of work.

Instead, AOL could use a centralized RADIUS server, as shown in Figure 1.3. Each AOL server is configured with a shared secret (similar to a password) and the RADIUS server is configured with a matching shared secret for each of the AOL servers.




Figure 1.3: Remote Authentication Dial-In User Service (RADIUS) configuration

This centralized RADIUS server would hold a centralized database of user accounts. Now if you tried to access AOL from any city, that AOL server would contact the RADIUS server to check your account.

Although this example with AOL works well to illustrate how RADIUS works, you don't need servers all over the country to take advantage of RADIUS. You could have as few as two or three RAS servers. Instead of having authentication databases on each server, you could configure a centralized RADIUS server.

RADIUS uses the User Datagram Protocol (UDP), which uses a best-effort delivery mechanism. In contrast, TACACS+ uses Transmission Control Protocol (TCP), which provides guaranteed delivery. In addition, RADIUS only encrypts the password, whereas TACACS+ encrypts the entire authentication process.

Diameter

Diameter is an extension of RADIUS and many organizations have switched to Diameter as a replacement for RADIUS due to its extra capabilities. One significant improvement is the support of the Extensible Authentication Protocol (EAP), which significantly enhances the security of Diameter. Additionally, Diameter adds several other commands beyond the capabilities of RADIUS. Diameter uses TCP instead of UDP used by RADIUS.

In geometry, the diameter of a circle is a straight line between the two edges of a circle, whereas the radius is a straight line from the center to an edge. In other words, the diameter of a circle is twice as long as the radius. The designers considered this when naming Diameter to indicate indirectly that it is twice as good as RADIUS.

Remember this

RADIUS provides centralized authentication. Diameter is an improvement over RADIUS, and it supports many additional capabilities, including securing transmissions with EAP.

XTACACS

Extended TACACS (XTACACS) is an older Cisco proprietary authentication protocol that is rarely used today. Most organizations use either RADIUS, Diameter, or TACACS+, instead of the older XTACACS.

TACACS+

Terminal Access Controller Access-Control System Plus (TACACS+) is the Cisco alternative to RADIUS, and is a recommended replacement for XTACACS. In addition to using TACACS+ for remote access, you can also use it for authentication with routers and other network devices.

TACACS+ provides two important security benefits over RADIUS. First, it encrypts the entire authentication process, whereas RADIUS encrypts only the password. Second, TACACS+ uses multiple challenges and responses between the client and the server.

Although TACACS+ is proprietary to Cisco, it can interact with Kerberos. This allows a Cisco RAS server (or VPN concentrator) to interact in a Microsoft Active Directory environment. As a reminder, Microsoft Active Directory uses Kerberos for authentication.

Organizations also use TACACS+ as an authentication service for network devices. In other words, you can use it to authenticate users before they are able to access a configuration page for a router or a switch. The network devices must be TACACS+ enabled, and a TACACS+ server provides the authentication services.

AAA Protocols

AAA protocols provide authentication, authorization, and accounting. Authentication verifies a user's identification. Authorization determines if a user should have access. Accounting tracks user access with logs.

As an example, RADIUS and TACACS+ are both considered AAA protocols because they provide all three services. They authenticate users who attempt remote access, determine if the user is authorized for remote access by checking a database, and then record the user's activity. TACACS+ uses multiple challenges and responses during a session. Kerberos is sometimes referred to as an AAA protocol, but it does not provide any accounting services.

Chapter 1 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Understanding Core Security Goals

- Confidentiality ensures that data is only viewable by authorized users. Encryption is the best choice to provide confidentiality. Access controls and steganography (hiding data inside other data) also protect the confidentiality of data.
- Integrity provides assurances that data has not been modified, tampered with, or corrupted through unauthorized or unintended changes. Data can be a message, a file, or data within a database. Hashing is a common method of ensuring integrity.
- Non-repudiation prevents entities from denying they took an action. Digital signatures and audit logs provide non-repudiation. Digital signatures also provide integrity for files and email.
- Availability ensures that data and services are available when needed. A common goal is to remove single points of failure. Methods used to increase or maintain availability include fault tolerance, failover clusters, load balancing, backups, virtualization, HVAC systems, and generators.
- Safety includes the safety of resources using physical security methods, such as fencing, lighting, door locks, and CCTV systems. It also includes the safety of personnel and can include escape plans, escape routes, and drills. Testing ensures these methods work as expected. Some electronic doors fail in an open state to ensure personnel safety.
- Layered security (or defense in depth) employs multiple layers of security to protect against threats. As an example, a firewall, an intrusion detection system, content filtering, and antivirus software provide multiple layers of protection. Security and IT professionals constantly monitor, update, add to, and improve existing security controls.

Introducing Basic Risk Concepts

- Risk is the possibility of a threat exploiting a vulnerability and resulting in a loss.
- A threat is any circumstance or event that has the potential to compromise confidentiality, integrity, or availability.
- A vulnerability is a weakness. It can be a weakness in the hardware, software, configuration, or users operating the system.
- Risk mitigation reduces risk by reducing the chances that a threat will exploit a vulnerability or by reducing the impact of the risk.
- Security controls reduce risks. For example, antivirus software is a security control that reduces the risk of malware infection.

Exploring Authentication Concepts

- Authentication allows entities to prove their identity by using credentials known to another entity.
- Identification occurs when a user claims or professes an identity, such as with a username, an email address, a PIV card, or by using biometrics.
- Authentication occurs when an entity provides proof of an identity (such as a password). A second identity is the authenticator and it verifies the authentication.
- Authorization provides access to resources based on a proven identity.
- Five factors of authentication are:
 - *Something you know*, such as a username and password
 - *Something you have*, such as a smart card, CAC, PIV, or a token
 - *Something you are*, using biometrics, such as fingerprints or retina scans
 - *Somewhere you are*, such as your location using geolocation technologies
 - *Something you do*, such as gestures on a touch screen
- The *something you know* factor typically refers to a shared secret, such as a password or a PIN. This is the least secure form of authentication.
- Passwords should be strong and changed often. Complex passwords include multiple character types. Strong passwords are complex and at least eight characters long.
- Administrators should verify a user's identity before resetting the user's password. When resetting passwords manually, administrators should configure them as temporary passwords that expire after the first use, requiring users to create a new password the first time they log on. Self-service password systems automate password recovery.
- Account lockout policies lock out an account after a user enters an incorrect password too many times.
- Smart cards are credit card-sized cards that have embedded certificates used for authentication. They require a PKI to issue certificates.
- Common Access Cards (CACs) and Personal Identity Verification (PIV) cards can be used as photo IDs and as smart cards (both identification and authentication).

- Tokens (or key fobs) display numbers in an LCD. These numbers provide rolling, one-time use passwords and are synchronized with a server. USB tokens include an embedded chip and a USB connection.
- HOTP and TOTP are open source standards used to create one-time-use passwords. HOTP creates a one-time-use password that does not expire and TOTP creates a one-time password that expires after 30 seconds.
- Biometric methods are the most difficult to falsify. Physical methods include fingerprints, retina scans, iris scans, and palm scans. Biometric methods can also be used for identification.
- Single-factor authentication includes one or more authentication methods in the same factor, such as a PIN and a password. Dual-factor (or two-factor) authentication uses two factors of authentication, such as a USB token and a PIN. Multifactor authentication uses two or more factors. Multifactor authentication is stronger than any form of single-factor authentication.
- Authentication methods using two or more methods in the same factor are single-factor authentication. For example, a password and a PIN are both in the *something you know* factor, so they only provide single-factor authentication.

Comparing Authentication Services

- Password Authentication Protocol (PAP) uses a password or PIN for authentication. A significant weakness is that PAP sends passwords across a network in cleartext.
- Challenge Handshake Authentication Protocol (CHAP) is more secure than PAP and uses a handshake process when authenticating clients. Both PAP and CHAP use PPP.
- Kerberos is a network authentication protocol using tickets issued by a KDC or TGT server. If a ticket-granting ticket expires, the user may not be able to access resources. Microsoft Active Directory domains and Unix realms use Kerberos for authentication.
- LDAP specifies formats and methods to query directories. It provides a single point of management for objects, such as users and computers, in an Active Directory domain or Unix realm. The following is an example of an LDAP string:
LDAP://CN=Homer,CN=Users,DC=GetCertifiedGetAhead,DC=com
- Secure LDAP encrypts transmissions with SSL or TLS.
- Single sign-on (SSO) allows users to authenticate with a single user account and access multiple resources on a network without authenticating again.
- SSO can be used to provide central authentication with a federated database and use this authentication in an environment with different operating systems (nonhomogeneous environment).
- SAML is an XML-based standard used to exchange authentication and authorization information between different parties. SAML is used with web-based applications.

Authenticating RAS Clients

- Remote access authentication is used when a user accesses a private network from a remote location, such as with a dial-up connection or a VPN connection.
- PAP sends passwords in cleartext.
- CHAP uses a challenge response authentication process.
- MS-CHAP and MS-CHAPv2 are the Microsoft improvement over CHAP. CHAPv2 provides mutual authentication.
- RADIUS provides central authentication for multiple remote access services. RADIUS relies on the use of shared secrets and only encrypts the password during the authentication process. It uses UDP.
- Diameter is an improvement over RADIUS and it supports EAP. Diameter uses TCP.
- XTACACS is a legacy protocol that is rarely used today.
- TACACS+ is used by some Cisco remote access systems as an alternative to RADIUS. TACACS+ uses TCP, encrypts the entire authentication process, and supports multiple challenges and responses.
- RADIUS, Diameter, and TACACS+ are all authentication, authorization, and accounting (AAA) protocols.

Chapter 1 Practice Questions

1. Homer needs to send an email to his HR department with an attachment that includes PII. He wants to maintain the confidentiality of this attachment. Which of the following choices is the BEST choice to meet his needs?
 - A. Hashing
 - B. Digital signature
 - C. Encryption
 - D. Certificate
2. You want to ensure that messages sent from administrators to managers arrive unchanged. Which security goal are you addressing?
 - A. Confidentiality
 - B. Integrity
 - C. Availability
 - D. Authentication
3. Your organization recently implemented two servers that act as failover devices for each other. Which security goal is your organization pursuing?
 - A. Safety
 - B. Integrity
 - C. Confidentiality
 - D. Availability
4. Management at your company recently decided to implement additional lighting and fencing around the property. Which security goal is your company MOST likely pursuing?
 - A. Confidentiality
 - B. Integrity
 - C. Availability
 - D. Safety
5. You are logging on to your bank's web site using your email address and a password. What is the purpose of the email address in this example?
 - A. Identification
 - B. Authentication
 - C. Authorization
 - D. Availability
6. Your organization has a password policy with a password history value of 12. What does this indicate?
 - A. Your password must be at least 12 characters long.
 - B. Twelve different passwords must be used before reusing the same password.
 - C. Passwords must be changed every 12 days.
 - D. Passwords cannot be changed until 12 days have passed.
7. A user calls into the help desk and asks the help-desk professional to reset his password. Which of the following choices is the BEST choice for what the help-desk professional should do before resetting the password?
 - A. Verify the user's original password.
 - B. Disable the user's account.
 - C. Verify the user's identity.
 - D. Enable the user's account.
8. Your organization is planning to implement remote access capabilities. Management wants strong authentication and wants to ensure that passwords expire after a predefined time interval. Which of the following choices BEST meets this requirement?
 - A. HOTP
 - B. TOTP

- C. CAC
- D. Kerberos

9. Which type of authentication is a fingerprint scan?

- A. Something you have
- B. Biometric
- C. PAP
- D. One-time password

10. When users log on to their computers, they are required to enter a username, a password, and a PIN. Which of the following choices BEST describes this?

- A. Single-factor authentication
- B. Two-factor authentication
- C. Multifactor authentication
- D. Mutual authentication

11. The security manager at your company recently updated the security policy. One of the changes requires dual-factor authentication. Which of the following will meet this requirement?

- A. Hardware token and PIN
- B. Fingerprint scan and retina scan
- C. Password and PIN
- D. Smart card

12. Your network infrastructure requires users to authenticate with something they are and something they know. Which of the following choices BEST describes this authentication method?

- A. Passwords
- B. Dual-factor
- C. Biometrics
- D. Diameter

13. Which of the following authentication services uses tickets for user credentials?

- A. RADIUS
- B. Diameter
- C. Kerberos
- D. LDAP

14. A network includes a ticket-granting ticket server. Which of the following choices is the primary purpose of this server?

- A. Authentication
- B. Identification
- C. Authorization
- D. Access control

15. Your network uses an authentication service based on the X.500 specification. When encrypted, it uses TLS. Which authentication service is your network using?

- A. SAML
- B. Diameter
- C. Kerberos
- D. LDAP

16. When you log on to your online bank account, you are also able to access a partner’s credit card site, check-ordering services, and a mortgage site without entering your credentials again. What does this describe?

- A. SSO

- B. Same sign-on
- C. SAML
- D. Kerberos

17. Your organization recently made an agreement with third parties for the exchange of authentication and authorization information. The solution uses an XML-based open standard. Which of the following is the MOST likely solution being implemented?

- A. RADIUS
- B. Diameter
- C. TACACS+
- D. SAML

18. Which of the following provides authentication services and uses PPP?

- A. Diameter and biometrics
- B. Kerberos and LDAP
- C. SAML and SSO
- D. PAP and CHAP

19. Users in your organization access your network from remote locations. Currently, the remote access solution uses RADIUS. However, the organization wants to implement a stronger authentication service that supports EAP. Which of the following choices BEST meets this goal?

- A. TACACS+
- B. Diameter
- C. Kerberos
- D. Secure LDAP

20. Which of the following choices provide authentication services for remote users and devices? (Select TWO.)

- A. Kerberos
- B. RADIUS
- C. Secure LDAP
- D. Diameter

Chapter 1 Practice Question Answers

1. **C.** Encryption is the best choice to provide confidentiality of any type of information, including Personally Identifiable Information (PII). Hashing, digital signatures, and certificates all provide integrity, not confidentiality.

2. **B.** Integrity provides assurances that data has not been modified, and integrity is commonly enforced with hashing. Confidentiality prevents unauthorized disclosure of data but doesn't address modifications of data. Availability ensures systems are up and operational when needed and uses fault tolerance and redundancy methods. Authentication provides proof that users are who they claim to be.

3. **D.** Your organization is pursuing availability. A failover cluster uses redundant servers to ensure a service will continue to operate even if one of the servers fail. Safety methods provide safety for personnel and other assets. Integrity methods ensure that data has not been modified. Confidentiality methods such as encryption prevent the unauthorized disclosure of data.

4. **D.** Lighting and fencing are two methods that can enhance the security goal of safety. Confidentiality is enhanced with encryption and access controls. Integrity is enhanced with hashing, certificates, and digital signatures. Availability is enhanced with redundancy and fault-tolerance procedures.

5. **A.** The email address provides identification for you and your account. The password combined with the email address provides authentication, proving who you are. Based on your identity, you are granted authorization to view your account details. Availability is unrelated to identification, authentication, and authorization.

6. **B.** The password history indicates how many passwords a system remembers and how many different passwords must be used before a password can be reused. Password length identifies the minimum number of characters. Password maximum age identifies when users must change passwords. Password minimum age identifies the length of time that must pass before users can change a password again.

7. **C.** Before resetting a user's password, it's important to verify the user's identity. Users often need the password reset because they have forgotten their original password, so it's not possible to verify the user's original password. It's not necessary to disable a user account to reset the password. You would enable the account if it was disabled or locked out, but the scenario doesn't indicate this is the case.

8. **B.** A Time-based One-Time Password (TOTP) meets this requirement. Passwords created with TOTP expire after 30 seconds. HMAC-based One-Time Password (HOTP) creates passwords that do not expire. A Common Access Card (CAC) is a type of smart card, but it does not create passwords. Kerberos uses tickets instead of passwords.
9. **B.** A fingerprint scan is a biometric method of authentication in the something you are factor of authentication. The something you have factor of authentication refers to something you can hold, such as a hardware token for a one-time password. Password Authentication Protocol (PAP) is an authentication method that sends passwords across the network in cleartext.
10. **A.** Both the password and the PIN are in the something you know factor of authentication, so this is single-factor authentication. Two-factor authentication requires the use of two different authentication factors. Multifactor authentication requires two or more factors of authentication. Mutual authentication is when both entities in the authentication process authenticate with each other and it doesn't apply in this situation.
11. **A.** A hardware token (such as an RSA token or a USB token) is in the something you have factor of authentication and the PIN is in the something you know factor of authentication. Combined, they provide dual-factor authentication. The remaining answers only provide single-factor authentication. A fingerprint scan and a retina scan are both in the something you are factor of authentication. A password and a PIN are both in the something you know factor of authentication. A smart card is in the something you have factor of authentication.
12. **B.** This is dual-factor authentication because users must authenticate with two different factors of authentication (something you are and something you know). Passwords are in the something you know factor and biometrics are in the something you are factor, but the scenario includes both factors, not just one. Diameter is a remote access authentication service that supports Extensible Authentication Protocol (EAP).
13. **C.** Kerberos uses a ticket-granting ticket server to create tickets for users and these tickets include user credentials for authentication. Remote Authentication Dial-In User Service (RADIUS) provides authentication for remote users. Diameter is an alternative to RADIUS and it can utilize Extensible Authentication Protocol (EAP). Lightweight Directory Access Protocol (LDAP) is an X.500-based authentication service.
14. **A.** Kerberos uses a ticket-granting ticket server for authentication. Users claim an identity with a username for identification. They prove their identity with credentials for authentication and Kerberos incorporates these credentials in tickets. Users are authorized access to resources with permissions, but only after they have been authenticated by an authentication service such as Kerberos. Access controls restrict access to resources after users are identified and authenticated.
15. **D.** Lightweight Directory Access Protocol (LDAP) uses X.500-based phrases to identify components and Secure LDAP can be encrypted with Transport Layer Security (TLS). Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML) used for single sign-on (SSO), but it is not based on X.500. Diameter is an alternative to Remote Authentication Dial-In User Service (RADIUS), but neither of these are based on X.500.
16. **A.** This is an example of single sign-on (SSO) capabilities because you can log on once and access all the resources without entering your credentials again. Same sign-on requires you to reenter your credentials for each new site, but you use the same credentials. Security Assertion Markup Language (SAML) is an SSO solution used for web-based applications and the bank might be using SAML, but other SSO solutions are also available. Kerberos is used in an internal network.
17. **D.** Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML) used for single sign-on (SSO) solutions. Remote Authentication Dial-In User Service (RADIUS) is a remote access authentication service. Diameter is an alternative to RADIUS. Terminal Access Controller Access-Control System Plus (TACACS+) is an authentication service that replaces the older TACACS protocol. RADIUS, Diameter, and TACACS+ do not use XML.
18. **D.** Both Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) use Point-to-Point Protocol (PPP). Diameter is an authentication service, but biometrics is an authentication method. Kerberos is an authentication service, but it doesn't use PPP and Lightweight Directory Access Protocol (LDAP) as a method of querying directories. Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML)-based data format used for single sign-on (SSO), but it doesn't use PPP.
19. **B.** Diameter is an alternative to Remote Authentication Dial-In User Service (RADIUS) and it can utilize Extensible Authentication Protocol (EAP). Terminal Access Controller Access-Control System Plus (TACACS+) is an authentication service that replaces older TACACS. Kerberos is an internal authentication protocol that uses tickets. Secure Lightweight Directory Access Protocol (LDAP) is an X.500-based authentication service that can be secured with Transport Layer Security (TLS).
20. **B, D.** Both Remote Authentication Dial-In User Service (RADIUS) and Diameter are authentication services for remote users and devices. Diameter is more secure than RADIUS. Kerberos is an authentication service used with a domain or realm and Secure Lightweight Directory Access Protocol (LDAP) uses Transport Layer Security (TLS) for encryption and is used to query directories.

Chapter 2

Exploring Control Types and Methods

CompTIA Security+ objectives covered in this chapter:

- 2.1 Explain the importance of risk related concepts.**
 - Control types (Technical, Management, Operational)
 - Importance of policies in reducing risk (Least privilege)
 - 2.3 Given a scenario, implement appropriate risk mitigation strategies.**
 - Perform routine audits
 - 2.6 Explain the importance of security related awareness and training.**
 - User habits (Prevent tailgating)
 - 2.7 Compare and contrast physical security and environmental controls.**
 - Physical security (Hardware locks, Mantraps, Video surveillance, Fencing, Proximity readers, Access list, Proper lighting, Signs, Guards, Barricades, Biometrics, Alarms, Motion detection)
 - Control types (Deterrent, Preventive, Detective, Compensating, Technical, Administrative)
 - 2.9 Given a scenario, select the appropriate control to meet the goals of security.**
 - Safety (Fencing, Lighting, Locks, CCTV)
 - 3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques.**
 - Detection controls vs. prevention controls (Camera vs. guard)
 - 4.3 Given a scenario, select the appropriate solution to establish host security.**
 - Hardware security (Cable locks, Safe, Locking cabinets)
 - 5.2 Given a scenario, select the appropriate authentication, authorization, or access control.**
 - Authorization (Least privilege, Mandatory access, Discretionary access, Rule-based access control, Role-based access control, Time of day restrictions)
 - Authentication (Access control)
 - 5.3 Install and configure security controls when performing account management, based on best practices.**
 - Account policy enforcement (Credential management, Group policy, Password complexity, Expiration, Disablement)
 - Group based privileges
 - User assigned privileges
- ****

Once you’ve ensured personnel have adequately identified themselves with authentication as described in Chapter 1, “Mastering Security Basics,” you can now move to different methods to restrict or control access. This chapter introduces basic security control types, and then covers common physical security and logical access controls. It ends with a description of four common access control models.

Understanding Control Types

Chapter 1 introduced the definitions of risk and risk mitigation. As a reminder, *risk* is the likelihood that a threat will exploit a vulnerability, resulting in a loss, and *risk mitigation* uses controls to reduce risk. You might see controls referred to as countermeasures or safeguards, referencing their ability to counter threats and provide safeguards to reduce vulnerabilities.

A *security incident* is an adverse event or series of events that can negatively affect the confidentiality, integrity, or availability of an organization's information technology (IT) systems and data. This includes intentional attacks, malicious software (malware) infections, accidental data loss, and much more. Security controls attempt to prevent or limit the impact of a security incident.

Just as there are many different types of security incidents, there are also many different types of security controls. You don't need to be an expert on all of them to pass the CompTIA Security+ exam, but you do need to have a basic understanding of some common controls. You also need to understand the three common implementation methods—technical, management, and operational—described in the following sections.

Control Implementation Methods

One method of classifying security controls is based on how they are implemented. The three common implementation classifications are technical, management, and operational:

- Technical controls use technology.
- Management controls use administrative or management methods.
- Operational controls are implemented by people in day-to-day operations.

Remember this

Security controls are classified as technical (implemented with technology), management (using administrative methods), and operational (for day-to-day operations).

Technical Controls

A *technical control* is one that uses technology to reduce vulnerabilities. An administrator installs and configures a technical control, and the technical control then provides the protection automatically. Throughout this book, you'll come across several examples of technical controls. The following list provides a few examples:

- **Encryption.** Encryption is a strong technical control used to protect the confidentiality of data. This includes data transferred over a network and data stored on devices such as servers, desktop computers, and mobile devices.
- **Antivirus software.** Once installed, the antivirus software provides protection against malware infection. Chapter 6, "Understanding Malware and Social Engineering," covers malware and antivirus software in more depth.
- **Intrusion detection systems (IDSs).** An IDS can monitor a network or host for intrusions and provide ongoing protection against various threats. Chapter 4, "Securing Your Network," covers different types of IDSs.
- **Firewalls.** Network firewalls restrict network traffic going in and out of a network. Chapter 3, "Understanding Basic Network Security," covers firewalls in more depth.
- **Least privilege.** The principle of least privilege specifies that individuals or processes are granted only the privileges they need to perform their assigned tasks or functions, but no more. Privileges are a combination of rights and permissions. The "Implementing Logical Access Controls" section later in this chapter covers least privilege in more depth.

The CompTIA Security+ exam focuses on many physical security and environmental controls. However, it's important to realize that many of these are also technical controls. For example, a security system that can detect motion and raise an alarm without user intervention is a technical control designed to increase physical security. Similarly, fire suppression systems use technologies to detect fires, raise alarms, and take various other actions to contain or extinguish the fires, all without user intervention. Fire suppression systems are environmental technical controls.

Remember this

Technical controls use technology to reduce vulnerabilities. Some examples include encryption, antivirus software, IDSs, firewalls, and the principle of least privilege. Technical physical security and environmental controls include motion detectors and fire suppression systems.

Management Controls

Management controls use planning and assessment methods to reduce and manage risk. Many provide an ongoing review of an organization's risk management capabilities. Some documents refer to management controls as administrative controls. Some common management controls are:

- **Risk assessments.** These help quantify and qualify risks within an organization so that the organization can focus on the serious risks. For example, a quantitative risk assessment uses cost and asset values to quantify risks based on monetary values. A qualitative risk assessment uses judgments to categorize risks based on probability and impact.
- **Vulnerability assessments.** A vulnerability assessment attempts to discover current vulnerabilities or weaknesses. When necessary, an organization implements additional controls to reduce the risk from these vulnerabilities.
- **Penetration tests.** These go a step further than a vulnerability assessment by attempting to exploit vulnerabilities. For example, a vulnerability assessment might discover a server isn't kept up to date with current patches, making it vulnerable to some attacks. A penetration test would attempt to compromise the server by exploiting one or more of the unpatched vulnerabilities.

Chapter 8, "Managing Risk," covers these assessments and tests in more depth. Some management controls focus on physical security and the environment. For example, an access list identifies individuals allowed into a secured area. Guards verify individuals are on the access list before allowing them in.

Operational Controls

Operational controls help ensure that day-to-day operations of an organization comply with their overall security plan. People (not technology) implement these controls. Operational controls include the following families:

- **Awareness and training.** The importance of training to reduce risks cannot be overstated. Training helps users maintain password security, follow a clean desk policy, understand threats such as phishing and malware, and much more.
- **Configuration and change management.** Configuration management often uses baselines to ensure that systems start in a secure, hardened state. Change management helps ensure that changes don't result in unintended configuration errors. Chapter 5, "Securing Hosts and Data," covers configuration and change management in more detail.
- **Contingency planning.** Chapter 9, "Preparing for Business Continuity," presents several different methods that help an organization plan and prepare for potential system outages. The goal is to reduce the overall impact on the organization if an outage occurs.
- **Media protection.** Media includes physical media such as USB flash drives, external and internal drives, and backup tapes.
- **Physical and environmental protection.** This includes physical controls, such as cameras, door locks, and environmental controls such as heating and ventilation systems.

NIST

The National Institute of Standards and Technology (NIST) is a part of the U.S. Department of Commerce, and it includes a Computer Security Division hosting the Information Technology Laboratory (ITL). The ITL publishes Special Publications (SPs) in the 800 series that are of general interest to the computer security community.

Many IT security professionals use these documents as references to design secure IT systems and networks. Additionally, many security-related certifications (beyond the CompTIA Security+ certification) also reference the SP 800 documents both directly and indirectly.

SP 800-53 Revision 4, “Security and Privacy Controls for Federal Information Systems and Organizations,” includes a wealth of information on security controls. It includes three relatively short chapters introducing security controls followed by multiple appendixes. Appendix F is a security control catalog that provides details on over 200 individual security controls, divided into 21 different families. For example, the Access Control family includes details on controls such as least privilege, account management, and separation of duties.

It’s worth noting that Revision 3 attempted to identify every control as technical, management, or operational. However, many controls included characteristics from more than just one of these classifications. NIST removed these references in Revision 4. That’s not to say that technical, management, and operational controls aren’t valid classifications. They are. However, it’s just not feasible to identify every control using only one classification.

If you’re interested in pursuing other security-related certifications or making IT security a career, the SP 800 documents are well worth your time. You can download SP 800-53 Revision 4 and other SP 800 documents at <http://csrc.nist.gov/publications/PubsSPs.html>.

. . .

Control Goals

Another way of classifying security controls is based on their goals in relationship to security incidents. Some common classifications are preventive, detective, corrective, deterrent, and compensating. The following sections describe these in more depth and the following list provides you with a short introduction:

- *Preventive controls* attempt to prevent an incident from occurring.
- *Detective controls* attempt to detect incidents after they have occurred.
- *Corrective controls* attempt to reverse the impact of an incident.
- *Deterrent controls* attempt to discourage individuals from causing an incident.
- *Compensating controls* are alternative controls used when a primary control is not feasible.

Preventive Controls

Ideally, an organization won't have any security incidents and that is the primary goal of preventive controls—to prevent security incidents. Some examples include:

- **Hardening.** Hardening is the practice of making a system or application more secure than its default configuration. This includes disabling unneeded services and protocols, protecting management interfaces and applications, protecting passwords, and disabling unnecessary accounts. Chapter 5 covers hardening in much more depth.
- **Security awareness and training.** Ensuring that users are aware of security vulnerabilities and threats helps prevent incidents. When users understand how social engineers operate, they are less likely to be tricked. For example, uneducated users might be tricked into giving a social engineer their passwords, but educated users will see through the tactics and keep their passwords secure.
- **Security guards.** Guards prevent and deter many attacks. For example, guards can prevent unauthorized access into secure areas of a building by first verifying user identities. Although attackers might attempt social engineering tactics to fool a receptionist, the presence of a guard will deter many social engineers from even trying these tactics.
- **Change management.** Change management ensures that changes don't result in unintended outages. In other words, instead of administrators making changes on the fly, they submit the change to a change management process. Notice that change management is an operational control, which attempts to prevent incidents. In other words, it's both an operational control and a preventive control.
- **Account disablement policy.** An account disablement policy ensures that user accounts are disabled when an employee leaves. This prevents anyone, including ex-employees, from continuing to use these accounts. The “Managing Accounts” section later in this chapter covers account disablement policies in more depth.

Remember this

Preventive controls attempt to prevent security incidents. Hardening systems increases a system's basic configuration to prevent incidents. Security guards can prevent unauthorized personnel from entering a secure area. Change management processes help prevent outages from configuration changes. An account disablement policy ensures that accounts are disabled when a user leaves the organization.

Detective Controls

Although preventive controls attempt to prevent security incidents, some will still occur. *Detective controls* attempt to detect when vulnerabilities have been exploited, resulting in a security incident. An important point is that detective controls discover the event after it's occurred. Some examples of detective controls are:

- **Log monitoring.** Several different logs record details of activity on systems and networks. For example, firewall logs record details of all traffic that the firewall blocked. By monitoring these logs, it's possible to detect incidents. Some automated methods of log monitoring automatically detect potential incidents and report them right after they've occurred.
- **Trend analysis.** In addition to monitoring logs to detect any single incident, you can also monitor logs to detect trends. For example, an intrusion detection system (IDS) attempts to detect attacks and raise alerts or alarms. By analyzing past alerts, you can identify trends such as an increase of attacks on a specific system.
- **Security audit.** Security audits can examine the security posture of an organization. For example, a password audit can determine if the password policy is ensuring the use of strong passwords. Similarly, a periodic review of user rights can detect if users have more permissions than they should.
- **Video surveillance.** A closed-circuit television (CCTV) system can record activity and detect what occurred. It's worth noting that video surveillance can also be used as a deterrent control.
- **Motion detection.** Many alarm systems are able to detect motion from potential intruders, and raise alarms.

Remember this

Detective controls attempt to detect when vulnerabilities have been exploited. Some examples include log monitoring, trend analysis, security audits, and CCTV systems.

Comparing Detection and Prevention Controls

It's worth stressing the differences between detection and prevention controls. A detective control can't predict when an incident will occur and it can't prevent it. In contrast, prevention controls stop the incident from occurring at all. Consider cameras and guards:

- **Video surveillance.** A simple camera without recording capabilities can prevent incidents because it acts as a deterrent. Compare this with a CCTV system with recording abilities. It includes cameras, which can deter and prevent incidents, but the full system is also a detection control because of the recording capabilities. Security professionals can review the recordings to detect incidents after they've occurred.
- **Guards.** Guards are primarily prevention controls. They will deter many incidents just by their presence. If attackers try to circumvent a security system, such as trying to sneak into a secured area, guards can intervene and stop them.

Corrective Controls

Corrective controls attempt to reverse the impact of an incident or problem after it has occurred. Some examples of corrective controls are:

- **Active IDS.** Active intrusion detection systems (IDSs) attempt to detect attacks and then modify the environment to block the attack from continuing. Chapter 4 covers IDSs in more depth.
- **Backups and system recovery.** Backups ensure that personnel can recover data if it is lost or corrupted. Similarly, system recovery procedures ensure administrators can recover a system after a failure. Chapter 9 covers backups and disaster recovery plans in more depth.

Deterrent Controls

Deterrent controls attempt to discourage a threat. Some deterrent controls attempt to discourage potential attackers from attacking, and others attempt to discourage employees from violating a security policy.

You can often describe many deterrent controls as preventive controls. For example, you can have a security guard control access to a restricted area of your building. This guard will deter most people from trying to sneak in. This deterrence prevents security incidents related to unauthorized access. Similarly, a social engineer might try to trick a building receptionist, but if you require visitors to go through the security guard first, it will deter many social engineers and prevent unauthorized entry.

The following list identifies some physical security controls used to deter threats:

- **Cable locks.** Securing laptops to furniture with a cable lock deters thieves from stealing the laptops. Thieves can't easily steal a laptop secured this way. If they try to remove the lock, they destroy the laptop. Admittedly, a thief could cut the cable with a large cable cutter. However, someone walking around with a four-foot cable cutter looks suspicious. If you're not familiar with a cable lock, refer to the "Using Hardware Locks" section later in this chapter.
- **Hardware locks.** Other locks such as locked doors securing a wiring closet or a server room also deter attacks. Many server bay cabinets also include locking cabinet doors.

Compensating Controls

Compensating controls are alternative controls used instead of a primary control. As an example, an organization might require smart cards as part of a multifactor authentication solution. However, it might take time for new employees to receive their smart card. To allow new employees to access the network and still maintain a high level of security, the organization might choose to implement a Time-based One-Time Password (TOTP) as a compensating control. The compensating control still provides multifactor authentication.

Combining Control Types and Goals

It's important to realize that the control types (technical, management, and operational) and control goals (preventive, detective, corrective, deterrent, and compensating) are not mutually exclusive. In other words, you can describe most controls with both terms.

As an example, encryption is a preventive technical control. It prevents the loss of data confidentiality, so it is a preventive control, and you implement it with technology, so it is a technical control. However, if you understand each of the types and each of the goals, you shouldn't have any problems picking out the correct answers on the exam, even if CompTIA combines them, such as a preventive technical control.

Comparing Physical Security Controls

A *physical security control* is something you can physically touch, such as a hardware lock, a fence, an identification badge, and a security camera.

Physical security access controls attempt to control entry and exits, and organizations commonly implement different controls at different boundaries, such as the following:

- **Perimeter.** Military bases and many other organizations erect a fence around the entire perimeter of their land. They often post security guards at gates to control access. In some cases, organizations install barricades to block vehicles.
- **Building.** Buildings commonly have additional controls for both safety and security. For example, guards and locked doors restrict entry so only authorized personnel enter. Many buildings include lighting and video cameras to monitor the entrances and exits.
- **Secure work areas.** Some companies restrict access to specific work areas when employees perform classified or restricted access tasks. In some cases, an organization restricts access to all internal work areas. In other words, visitors can enter the lobby of a building, but they are not able to enter internal work areas without an escort.
- **Server and network rooms.** Servers and network devices such as routers and switches are normally stored in areas where only the appropriate IT personnel can access them. These spaces may be designated as server rooms or wiring closets. It's common for an organization to provide additional physical security for these rooms to prevent attackers from accessing the equipment. For example, locking a wiring closet prevents an attacker from installing illicit monitoring hardware, such as a protocol analyzer, to capture network traffic.
- **Hardware.** Additional physical security controls protect individual systems. For example, server rooms often have locking cabinets to protect servers and other equipment installed in the equipment bays. Cable locks protect laptop computers, and smaller devices can be stored in safes.

Comparing Door Access Systems

A *door access system* is one that only opens after some access control mechanism is used. Some common door access controls are cipher locks, proximity cards, and biometrics.

When implementing door access systems, it's important to limit the number of entry and exit points. As an example, if a data center has only one entrance and exit, it is much easier to monitor this single access point. You can control it with door locks, video surveillance, and guards. On the other hand, if the data center has two entry/exit points, you need another set of controls to control access in both places.

Another important consideration with door access systems is related to personnel safety and fire. In the event of a fire, door access systems should allow personnel to exit the building without any form of authentication.

Remember this

In the event of a fire, door access systems should allow personnel to exit the building without any form of authentication. Access points to data centers and server rooms should be limited to a single entrance and exit whenever possible.

Securing Door Access with Cipher Locks

Cipher locks often have four or five buttons labeled with numbers. Employees press the numbers in a certain order to unlock the door. For example, the cipher code could be 1, 3, 2, 4. Users enter the code in the correct order to gain access. Cipher locks can be electronic or manual. An electronic cipher lock automatically unlocks the door after you enter the correct code into the keypad. A manual cipher lock requires the user to turn a handle after entering the code.

To add complexity and reduce brute force attacks, many manual cipher locks include a code that requires two numbers entered at the same time. Instead of just 1, 3, 2, 4, the code could be 1/3 (entered at the same time), then 2, 4, 5.

One challenge with cipher locks is that they don't identify the users. Further, uneducated users can give out the cipher code to unauthorized individuals without understanding the risks. Shoulder surfers might attempt to discover the code by watching users as they enter it. Security awareness training can help reduce these risks.

Remember this

Cipher locks require users to enter a code to gain access. It's important to provide training to users on the importance of keeping the code secure. This includes not giving it out to others and preventing shoulder surfers from seeing the code when users enter it. Cipher locks do not identify users.

Securing Door Access with Proximity Cards

Proximity cards are small credit card-sized cards that activate when they are in close proximity to a card reader. Many organizations use these for access points, such as the entry to a building or the entry to a controlled area within a building. The door uses an electronic lock that only unlocks when the user passes the proximity card in front of a card reader.

You've probably seen proximity card readers implemented with credit card readers. Many self-serve gasoline stations and fast-food restaurants use them. Instead of swiping your credit card through a magnetic reader, you simply pass it in front of the reader (in close proximity to the reader), and the reader extracts your credit card's information.

It's intriguing how this is accomplished. The card doesn't require its own power source. Instead, the electronics in the card include a capacitor and a coil that can accept a charge from the proximity card reader. When you pass the card close to the reader, the reader excites the coil and stores a charge in the capacitor. Once charged, the card transmits the information to the reader using a radio frequency.

When used with door access systems, the proximity card can send just a simple signal to unlock the door. Some systems include details on the user and record when the user enters or exits the area. When used this way, it's common to combine the proximity card reader with a key pad requiring the user to enter a personal identification number (PIN). This identifies and authenticates the user with multifactor authentication. The user has something (the proximity card) and knows something (a PIN).

Many organizations use proximity cards with turnstiles to provide access for a single person at a time. These are the same type of turnstiles used as entry gates in subways, stadiums, and amusement parks.

Remember this

Proximity cards are credit card-sized access cards. Users pass the card near a proximity card reader and the card reader then reads data on the card. Some access control points use proximity cards with PINs for authentication.

As a side note, one of the worrisome issues related to these cards is that attackers can build or purchase systems that can read your credit cards if they operate as proximity cards. The attacker places the reader in a purse or bag and positions it close to your wallet or purse, perhaps by standing

behind you in the elevator, a store, or a line. The electronics on the card would charge and then transmit without your knowledge. The collected information can be used later to make unauthorized purchases.

About the only way to prevent this is to wrap your credit cards in some type of shielding, like aluminum foil. There are companies that sell credit card shield protectors for as much as \$29.95. However, you can make your own shield with a couple of well-placed pieces of aluminum foil in your wallet or purse.

Securing Door Access with Biometrics

It’s also possible to use biometric methods as an access control system. One of the benefits is that some biometric methods provide both identification and authentication. When connected to a back-end database, these systems can easily record the activity, such as who entered the area and when.

For example, you can install a retina scanner at the entrance to a secure server room. When individuals want to enter, the biometric scanner identifies and authenticates them. It’s important to ensure you use a biometric system with a low false acceptance rate, as described in Chapter 1. Otherwise, it might falsely identify unauthorized individuals and grant them access.

Remember this

Door access systems include cipher locks, proximity cards, and biometrics. Cipher locks do not identify users. Proximity cards can identify and authenticate users when combined with a PIN. Biometrics can also identify and authenticate users.

Identifying Users with ID Badges

An *identification badge* (ID badge) identifies its owner with a picture and other related information such as a name. Many organizations require employees to wear ID badges while at work. Employees clip the badge to a pocket or shirt, or connect it to a lanyard that they wear around their neck. Anyone without an ID badge easily sticks out as an unauthorized intruder.

You can also combine other electronic measures with an ID badge. For example, as mentioned in Chapter 1, government agencies use Common Access Cards (CACs) and Personal Identity Verification (PIV) cards for both identification and as a smart card. Similarly, some organizations include proximity card technology on the ID badge so that employees can use their badges to open doors. Other organizations use magnetic strips on the card, which includes information on the user.

Tailgating

Tailgating (also called piggybacking) occurs when one user follows closely behind another user without using credentials. For example, if Lisa opens a door with her proximity card and Bart follows closely behind her without using a proximity card, Bart is tailgating. If authorized users routinely do this, it indicates the environment is susceptible to a social engineering attack where an unauthorized user follows closely behind an authorized user.

As an example, an organization hired a security company to perform a vulnerability assessment. The company sent one of its top security professionals (who happened to be an attractive woman) to see if she could get into the building. She saw that employees were using proximity cards to get into the building, but she didn't have one. Instead, she loaded herself up with a book bag and a laptop—ensuring her hands weren't free. She timed her approach carefully and followed closely behind an employee with a proximity card. She flashed a friendly smile, and sure enough, the employee held the door open for her.

Most of us learn to be polite and courteous and social engineers take advantage of this. It's polite to hold a door open for people who have their hands full. In contrast, it's rude to slam the door in the face of someone following behind us. However, most users don't want to help criminals. Security awareness programs and training help users understand how criminals use tactics such as tailgating. Educated users are less likely to be tricked, even by a friendly smile from an attractive woman.

High-traffic areas are most susceptible to tailgating attacks. Security guards can be an effective preventive measure at access points, but they need to be vigilant to ensure that tailgating does not occur. The best solution is a mantrap.

Preventing Tailgating with Mantraps

A *mantrap* is a physical security mechanism designed to control access to a secure area through a buffer zone. Personnel use something like a proximity card to gain access, and the mantrap allows one person, and only one person, to pass through. Because they only allow one person through at a time, mantraps prevent tailgating. Mantraps get their name due to their ability to lock a person between two areas, such as an open access area and a secure access area, but not all of them are that sophisticated.

An example of a simple mantrap is a turnstile similar to what you see in many public transport systems. Even if you've never ridden the subway in one of many U.S. cities or the Tube in London, you've probably seen turnstiles in movies such as *While You Were Sleeping*. When customers present a token, the turnstile unlocks and allows a single person through at a time. Similarly, users unlock the turnstile mantrap with something like a proximity card.

A sophisticated mantrap is a room, or even a building, that creates a large buffer area between the secure area and the unsecured area. Access through the entry door and the exit door is tightly controlled, either with guards or with an access card such as a proximity card.

It's also possible to require identification and authentication before allowing passage through a mantrap. For example, a retina scanner can identify individuals and restrict access to only authorized individuals. Similarly, some card reader systems support the use of unique PINs assigned to the user. Users present their card and enter their PIN to gain access before the mantrap opens.

Remember this

Tailgating is a security violation that occurs when one user follows closely behind another user without using credentials. Mantraps allow only a single person to pass at a time. Sophisticated mantraps can identify and authenticate individuals before allowing access.

Increasing Physical Security with Guards

Many organizations use security guards to control access to buildings and secure spaces. If employees have ID badges, guards can check these badges prior to granting the employees access. Even if ID badges aren't used, guards can still verify people's identity using other identification. Similarly, the security guards can restrict access by checking people's identity against a preapproved access control list.

Security guards can also take a less-active role to deter security incidents. For example, a security guard can deter tailgating incidents by observing personnel when they use their proximity card to gain access to a secure area.

Remember this

Security guards are physical security controls that can protect access to restricted areas. Security guards can be an effective deterrent to prevent tailgating. They can also check individuals' identification against a preapproved access list.

Controlling Access with Access Lists and Logs

Some organizations maintain physical access lists identifying authorized personnel. Guards staff entry points, and check the access list prior to letting people in. If people aren't on the list, the guard doesn't let them in. Additionally, the guards can maintain physical access logs to document when employees enter or exit.

When access points include electronic measures such as proximity card readers that record information on users, electronic access logs can verify exactly when someone enters or exits a building.

You can also use these logs to identify security vulnerabilities. For example, if a log shows that a user exited a building, but does not show that the user entered the building, it indicates the user probably tailgated when entering.

Although these access logs are very useful, it's worth noting that video surveillance provides the most reliable proof of a person's location and activity. People can steal a proximity card, but they can't steal a person's likeness shown on a video.

Monitoring Areas with Video Surveillance

Organizations are increasingly using security cameras in the workplace and surrounding areas for video surveillance. This includes areas outside of a building, such as a parking lot, and all building entrances and exits. Additionally, many organizations use video surveillance to monitor internal entrances of high-security areas, such as the entrance of a data center or server room.

Video surveillance is often referred to as closed-circuit television (CCTV) because it transmits signals from video cameras to monitors that are similar to TVs. In addition to providing security, CCTV can also enhance safety by deterring threats.

Organizations often use video cameras within a work environment to protect employees and enhance security in the workplace. In addition to live monitoring, most systems include a recording element, and they can verify if someone is stealing the company's assets. By recording activity, videos can be played back later for investigation and even prosecution.

Video surveillance provides the most reliable proof of a person's location and activity. Access logs provide a record, but it's possible to circumvent the security of an access log. For example, if Bart used your proximity card to gain access to a secure space, the log will indicate you entered, not Bart. In contrast, if the video shows that Bart entered the room at a certain time of day, it's not easy for Bart to refute the video.

Remember this

Video surveillance provides reliable proof of a person's location and activity. It can identify who enters and exits secure areas and can record theft of assets.

When using video surveillance in a work environment, it's important to respect privacy and to be aware of privacy laws. Some things to consider are:

- **Only record activity in public areas.** People have a reasonable expectation of privacy in certain areas, such as locker rooms and restrooms, and it is often illegal to record activity in these areas.
- **Notify employees of the surveillance.** If employees aren't notified of the surveillance, legal issues related to the video surveillance can arise. This is especially true if the recordings are used when taking legal and/or disciplinary actions against an employee.
- **Do not record audio.** Recording audio is illegal in many jurisdictions, without the express consent of all parties being recorded. Many companies won't even sell surveillance cameras that record audio.

Combining Fencing and Motion Detection

Fences provide a barrier around a property and deter people from entering. In addition to providing security, fences also enhance safety by preventing unauthorized individuals from entering a secure area.

You can increase the deterrent effects by using higher fences and barbed wire, though this isn't always possible. For example, a manufacturing facility surrounded by coils of barbed wire evokes memories of war and might give the company a bad reputation. Similarly, some building codes restrict the height of some fences.

An alternative to higher fences is the use of motion detection sensors. These detect any motion on the fence, such as when someone is trying to climb it. A challenge with these fences is the high incidence of false alarms. For example, a cat or squirrel climbing the fence might trigger the motion detector.

Combining Proper Lighting and Motion Detection

Installing lights at all of the entrances to a building can deter attackers from trying to break in. Similarly, lighting at the entrances of any internal restricted areas can deter people from trying to enter. In addition to enhancing security, lighting also contributes to personnel safety.

Many organizations use a combination of automation, light dimmers, and motion sensors to save on electricity costs without sacrificing security. The lights automatically turn on at dusk, but in a low, dimmed mode. When the motion sensors detect any movement, the lights turn on at full capacity. They automatically turn off at dawn.

It's important to protect the lights. For example, if an attacker can remove the light bulbs, it defeats the control. Either place the lights high enough so that they can't be easily reached, or protect them with a metal cage.

Combining Alarms and Motion Detection

Alarms provide an additional physical security protection. This includes alarms that detect fire and alarms that detect unauthorized access. Fire alarms detect smoke and/or heat and trigger fire suppression systems. Burglary prevention systems monitor entry points such as doors and windows, detecting when someone opens them.

You can also combine motion detection systems with burglary prevention systems. They detect movement within monitored areas and trigger alarms. Obviously, you wouldn't have motion detection systems turned on all the time. Instead, you'd turn them on when people will not be working in the area, such as during nights or weekends.

Securing Access with Barricades

In some situations, fencing isn't enough to deter potential attackers. To augment fences and other physical security measures, organizations erect stronger barricades. As an example, military bases often erect strong, zigzag barricades that require vehicles to slow down to navigate through them. This prevents attackers from trying to ram through the gates.

Businesses and organizations need to present an inviting appearance, so they can't use such drastic barricades. However, they often use bollards, which are short vertical posts, composed of reinforced concrete and/or steel. They often place the bollards in front of entrances, placed about three or four feet apart. They typically paint them with colors that match their store. You've probably walked through a set of bollards multiple times without giving them a second thought. However, thieves who are contemplating driving a car or truck through the entrance do see them.

Many thieves have driven vehicles right through the front of buildings, and then proceeded to steal everything in sight. Depending on the strength of the walls, criminals might even be able to drive through a wall with a truck. Strategically placed bollards will prevent these types of attacks.

Remember this

Barricades provide stronger barriers than fences and attempt to deter attackers. Bollards are effective barricades that can block vehicles.

Using Signs

A simple physical security control is a sign. For example, an “Authorized Personnel Only” sign will deter many people from entering a restricted area. Similarly, “No Trespassing” signs let people know they shouldn’t enter. Of course, these signs won’t deter everyone, so an organization typically uses additional physical security measures.

Using Hardware Locks

You can implement simple physical security measures to prevent access to secure areas. For example, you can use hardware locks—similar to what you use to secure your home—to secure buildings as well as rooms within buildings. Companies that don't have the resources to employ advanced security systems often use these types of hardware locks.

Instead of allowing free access to wiring closets or small server rooms, small organizations use these types of locks to restrict access. Although these locks aren't as sophisticated as the ones used by large organizations, they are much better than leaving the rooms open and the equipment exposed.

Securing Mobile Computers with Cable Locks

Cable locks are a great theft deterrent for mobile computers, and even many desktop computers at work. Computer cable locks work similar to how a bicycle cable lock works. However, instead of securing a bicycle to a bike rack or post, a computer cable lock secures a computer to a piece of furniture.

The user wraps the cable around a desk, table, or something heavy, and then plugs it into an opening in the laptop specifically created for this purpose. Most cable locks have a four-digit combo. If you (or anyone) remove the cable lock without the combo, it will likely destroy the laptop.

Another common use of cable locks is for computers in unsupervised labs. For example, you can secure laptop or desktop computers with cable locks in a training lab. This allows you to leave the room open so that students can use the equipment, but the cable locks prevent thieves from stealing the equipment.

Remember this

Cable locks are effective threat deterrents for small equipment such as laptops and some workstations. When used properly, they prevent losses due to theft of small equipment.

Securing Servers with Locking Cabinets

Larger companies often have large server rooms with advanced security to restrict access. Additionally, within the server room, administrators use locking cabinets to secure equipment mounted within the bays. An equipment bay is about the size of a large refrigerator and can hold servers, routers, and other IT equipment. These bays have doors in the back and many have doors in the front, too. Administrators lock these doors to prevent unauthorized personnel from accessing the equipment.

Offices often have file cabinets that lock, too, so it's important to pay attention to the context when referring to locking cabinets. For example, if you want to secure equipment within a server room, locking cabinets is one of many physical security controls you can use. If you want to secure unattended smartphones in an office space, you can use also use a locking cabinet, but this is an office file cabinet that locks.

Remember this

Locking cabinets in server rooms provide an added physical security measure. A locked cabinet prevents unauthorized access to equipment mounted in server bays.

Securing Small Devices with a Safe

Locking file cabinets or safes used in many offices help prevent the theft of smaller devices. For example, you can store smaller devices such as external USB drives or USB flash drives in an office safe or locking cabinet when they aren't in use. Depending on the size of the office safe and office cabinet, you might also be able to secure laptops within them.

Implementing Logical Access Controls

Logical access control methods are implemented through technologies such as Group Policy and account management tools. They control access to the logical network as opposed to controlling access to the physical areas of a building or physical access to devices within the network.

When studying these methods, it's valuable to understand some underlying principles, such as the principle of least privilege and the principle of need to know. This section covers these principles and then digs into the logical access controls.

Least Privilege

The principle of least privilege is an example of a technical control implemented with access controls. Privileges are the rights and permissions assigned to authorized users. Least privilege specifies that individuals and processes are granted only the rights and permissions needed to perform assigned tasks or functions, but no more. For example, if Lisa needs read access to a folder on a server, you should grant her read access to that folder, but nothing else.

A primary goal of implementing least privilege is to reduce risks. As an example, imagine that Carl works at the Nuclear Power Plant, but administrators have not implemented the principle of least privilege. In other words, Carl has access to all available data within the Nuclear Power Plant, not just the limited amount of data he needs to perform his job. Later, Lenny gets into trouble and needs money, so he convinces Carl to steal data from the power plant so that they can sell it. In this scenario, Carl can steal and sell all the data at the plant, which can result in serious losses.

In contrast, if administrators applied the principle of least privilege, Carl would only have access to a limited amount of data. Even if Lenny convinces him to steal the data, Carl wouldn't be able to steal very much simply because he doesn't have access to it. This limits the potential losses for the power plant.

This principle applies to regular users and administrators. As an example, if Marge administers all the computers in a training lab, it's appropriate to give her administrative control over all these computers. However, her privileges don't need to extend to the domain, so she wouldn't have administrative control over all the computers in a network. Additionally, she wouldn't have the privileges required to add these computers to the domain, unless that was a requirement in the training lab. Similarly, if a network administrator needs to review logs and update specific network devices, it's appropriate to give the administrator access to these logs and devices, but no more.

Many services and applications run under the context of a user account. These services have the privileges of this user account, so it's important to ensure that these accounts are only granted the privileges needed by the service or the application. In the past, many administrators configured these service and application accounts with full administrative privileges. When attackers compromised a service or application configured this way, they gained administrative privileges and wreaked havoc on the network.

Remember this

Least privilege is a technical control. It specifies that individuals or processes are granted only those rights and permissions needed to perform their assigned tasks or functions.

Need to Know

The principle of need to know is similar to the principle of least privilege in that users are granted access only to the data and information that they need to know for their job. Notice that need to know is focused on data and information, which is typically protected with permissions. In contrast, the principle of least privilege includes both rights and permissions.

Rights refer to actions and include actions such as the right to change the system time, the right to install an application, or the right to join a computer to a domain. Permissions typically refer to permissions on files, such as read, write, and modify.

Group Policy

Windows domains use Group Policy to manage multiple users and computers in a domain. Group Policy allows an administrator to configure a setting once in a Group Policy Object (GPO) and apply this setting to many users and computers within the domain. Although you can implement Group Policy on single, stand-alone Windows computers, the great strength of Group Policy comes when you implement it in a domain. Administrators implement domain Group Policy on domain controllers.

As an example, if you want to change the local Administrator password on all the computers in your domain, you can configure a GPO once, link the GPO to the domain, and it changes the local Administrator password for all the computers in the domain. The magic of Group Policy is that it doesn't matter if you have five systems or five thousand systems. The policy still only needs to be set once to apply to all of the systems.

Administrators also use Group Policy to target specific groups of users or computers. For example, Active Directory allows you to organize user accounts and computer accounts into organizational units (OUs). You can create a GPO and link it to a specific OU, and the GPO settings only apply to the users and computers within the OU. These settings do not apply to users and computers in other OUs.

Remember this

Group Policy is implemented on a domain controller within a domain. Administrators use it to create password policies, lock down the GUI, configure host-based firewalls, and much more.

Using a Password Policy

A common group of settings that administrators configure in Group Policy is the password policy settings. A password policy ensures that users create strong passwords and change them periodically.

Password policies typically start as a written document that identifies the organization's security goals related to passwords. For example, it might specify that passwords must be at least eight characters long, complex, and users should change them every 45 days. Administrators then implement these requirements with a technical control such as a technical password policy within a GPO.

Figure 2.1 shows the Local Group Policy Editor with the password policy selected in the left pane. The right pane shows the password policy for a Windows system and the following text explains these settings:

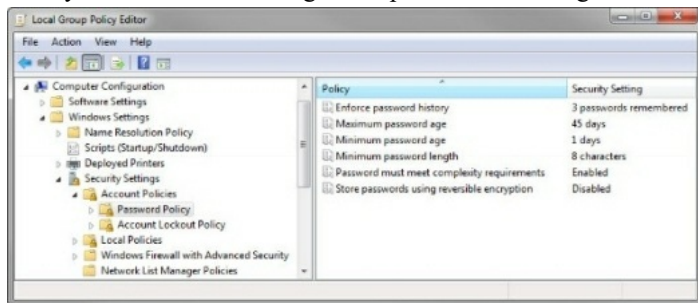


Figure 2.1: Password policy in a Windows domain

- **Enforce password history.** Some users will go back and forth between two passwords that they constantly use and reuse. However, password history remembers past passwords and prevents the user from reusing previously used passwords. For example, setting this to 3 prevents users from reusing passwords until they've used 3 new passwords. Similarly, setting it to 24 prevents users from reusing passwords until they've used 24 new passwords.
- **Maximum password age.** This setting defines when users must change their password. For example, setting this to 45 days causes the password to expire after 45 days. This forces users to reset their password to a new password on the 46th day.
- **Minimum password age.** The minimum password age defines how long users must wait before changing their password again. If you set this to 1 day, it prevents users from changing their passwords until 1 day has passed. This is useful with a password history to prevent users from changing their password multiple times until they get back to the original password. If the password history is set to 24 and the minimum password age is set to 1 day, it'll take a user 25 days to get back to the original password. This is enough to discourage most users.
- **Minimum password length.** This setting enforces the character length of the password. It's common to require users to have passwords at least 8 characters long, but some organizations require users to have passwords as long as 15 characters.
- **Password must meet complexity requirements.** This setting requires users to have complex passwords that include at least three of the four character types (uppercase letters, lowercase letters, numbers, and special characters). You might remember from Chapter 1 that you increase the key space (or complexity) of a password by using more character types.
- **Store passwords using reversible encryption.** Reversible encryption stores the password in such a way that the original password can be discovered. This is rarely enabled.

If an administrator creates the initial password for a user or resets the password, the administrator should set the password to expire immediately. This is the same as a password reaching its maximum age and it forces the user to reset the password immediately when logging on for the first time. If

more than one person knows the credentials for an account, the credentials no longer uniquely identify the individual—someone else could log on with the same credentials, and even if the system logged the event, you can’t prove who did it.

Remember this

Password policies include several elements. The password history is used with the minimum password age to prevent users from changing their password to a previously used password. Maximum password age causes passwords to expire and requires users to change their passwords periodically. Minimum password length specifies the minimum number of characters in the password. Password complexity increases the key space, or complexity, of a password by requiring more character types such as special characters.

Domain Password Policy

In Windows domains, a single password policy applies to all users in the domain. It doesn’t matter if the domain has five users or five thousand users. When an administrator sets the password policy, it applies to all users.

Although most domains have only a single password policy that applies to all user accounts, it is possible to create additional password policies for different groups. For example, many organizations want administrators to use strong passwords of at least 15 characters, but allow users to create passwords with only 8 characters.

To achieve this, security administrators start with a standard password policy that applies to all users in the domain. Next, they create a special password policy and apply it to specific groups of users such as specially created administrative accounts.

Application Passwords

Group Policy doesn’t necessarily apply to passwords used within applications. If the application runs under the context of a user account, then the password policy will apply to that application user account. For example, if an application needs a user account, an administrator can create a user account named AppAccount. Other than the name, it works just like any user account and the domain password policy applies to it.

However, some applications use internal passwords completely separate from Group Policy. In these situations, the application developers should still identify the organization’s password goals and ensure they develop the application to comply with the policy.

Remember this

Password policies should apply to any entity using passwords. This includes user accounts and any accounts used by services and applications. Applications using internally developed password policies should still adhere to an organization’s password policy.

Managing Accounts

Account management is concerned with the creation, management, disabling, and termination of accounts. When the account is active, access control methods are used to control what the user can do. Additionally, administrators use access controls to control when and where users can log on.

Disabling and Deleting Accounts

Many organizations have account management policies (sometimes called account disablement policies) that specify how to manage accounts in different situations. For example, most organizations require administrators to disable user accounts as soon as possible when employees leave the organization.

Disabling is preferred over deleting the account, at least initially. If administrators delete the account, they also delete any encryption and security keys associated with the account. However, these keys are retained when the account is disabled. As an example, imagine that an employee encrypted files with his account. If the account was deleted, these files may remain encrypted forever unless the organization has a key escrow or recovery agent that can access the files.

Some contents of an account disablement policy include:

- **Terminated employee.** An account disablement policy specifies that accounts for ex-employees are disabled as soon as possible. This ensures a terminated employee doesn't become a disgruntled ex-employee who wreaks havoc on the network. Note that "terminated" refers to both employees who resign and employees who are fired.
- **Leave of absence.** If an employee will be absent for an extended period, the account should be disabled while the employee is away. Organizations define *extended period* differently, with some organizations defining it as only two weeks, whereas other organizations extend it out to as long as two months.
- **Delete account.** When the organization determines the account is no longer needed, administrators delete it. For example, the policy may direct administrators to delete accounts that have been inactive for 60 or 90 days.

Many administrators routinely run scripts to identify inactive accounts. For example, it's relatively simple in a Microsoft domain to run a script listing all enabled accounts that haven't been used in the last 30 days. Administrators disable these accounts to prevent their use. Often these are accounts of ex-employees or temporary employees who are no longer at the organization. Ideally, the account disablement policy would ensure that the accounts are disabled as soon as the employee leaves, but the scripts provide an additional check to ensure inactive accounts are disabled.

Remember this

An account disablement policy identifies what to do with accounts for employees who leave permanently or on a leave of absence. Most policies require administrators to disable the account as soon as possible, so that ex-employees cannot use the account. Disabling the account ensures that data associated with it remains available. Security keys associated with an account remain available when the account is disabled, but are no longer accessible if the account is deleted.

Recovering Accounts

In some situations, administrators need to recover accounts. The two primary account recovery scenarios are:

- **Enable a disabled account.** Administrators can reset the user's password and take control of the account. Similarly, they pass control of the account to someone else such as a supervisor or manager of an ex-employee. Administrators reset the user's password, set it to expire on first use, and then give the password to the other person.
- **Recover a deleted account.** It is also possible to recover a deleted account. This is more complex than simply creating another account with the same name. Instead, administrators follow detailed procedures to recover the account.

Prohibiting Generic Accounts

Another element of an account management policy prohibits the use of generic accounts. In Microsoft operating systems and domains, this means that personnel should not use the Guest account. Instead, administrators ensure the Guest account remains disabled.

As a reminder, access control requires that the system identifies and authenticates each user before granting them access. Imagine that five people are using the Guest account. In this scenario, users are not identified. If one of the users deletes all the files shared by the five users, it's not possible to determine who actually deleted the files even if the system included detailed logging. The log would indicate that the files were deleted by the Guest account, but the actual person who deleted them could have been any of the five users.

Restricting Access Based on Time-of-Day

Time-of-day restrictions specify when users can log on to a computer. If a user tries to log on to the network outside the restricted time, the system denies access to the user.

As an example, imagine a company operates between 8:00 a.m. and 5:00 p.m. on a daily basis. Managers decide they don't want regular users logging on to the network except between 6:00 a.m. and 8:00 p.m., Monday through Friday. You could set time-of-day restrictions for user accounts, as shown in Figure 2.2. If a user tries to log on outside the restricted time (such as during the weekend), the system prevents the user from logging on.

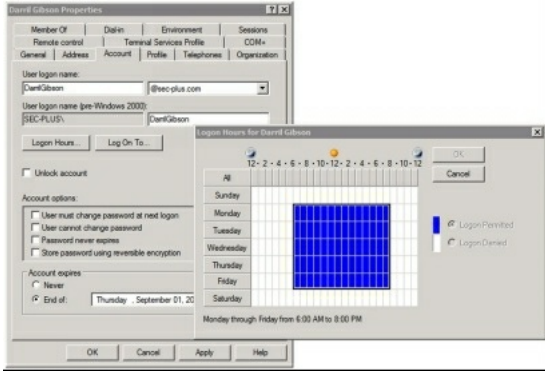


Figure 2.2: User account properties with time restrictions

If users are working overtime on a project, the system doesn't log them off when the restricted time arrives. For example, if Maggie is working late on a Wednesday night, the system doesn't log her off at 8:00 p.m. (assuming the time restrictions are set as shown in Figure 2.2). However, the system will prevent her from creating any new network connections.

Remember this

Time-of-day restrictions prevent users from logging on during restricted times. It also prevents logged-on users from accessing resources during certain times.

Expiring Accounts

It's possible to set user accounts to expire automatically. When the account expires, the system disables it, and the user is no longer able to log on using the account.

If you look back at Figure 2.2, it shows the properties of an account. The Account Expires section is at the bottom of the page, and the account is set to expire on September 1. When September 1 arrives, the account is automatically disabled and the user will no longer be able to log on.

It's common to configure temporary accounts to expire. For example, an organization may hire contractors for a 90-day period to perform a specific job. An administrator creates accounts for the contractors and sets them to expire in 90 days. This automatically disables the accounts at the end of the contract. If the organization extends the contract, it's a simple matter to change the expiration date and enable the account.

Remember this

Account expiration dates automatically disable accounts on the expiration date. This is useful for temporary accounts such as temporary contractors.

Reviewing Account Access

Configuring logging of logon attempts is an important security step for system monitoring. After configuring logging, a system records the time and date when users log on, and when they access systems within a network. When users first log on to their account, it's recorded as a logon action. Additionally, when users access a resource over the network (such as a file server), it is also recorded as a logon action. Many systems utilize single sign-on, so users don't have to provide their credentials again. However, their access is still recorded as a logon action.

You can identify if someone is trying to hack into an account by monitoring failed logon attempts. If a log shows 50 failed logon attempts followed by a success, it indicates someone successfully guessed the password for an account.

Chapter 1 presented information related to account lockout policies where a system locks out an account after so many failed logon attempts. However, the administrator account cannot be locked out. If the name of the administrator account is not changed (a standard security practice) and someone tries to hack into it, an account logon audit will capture the details.

Chapter 8 covers logs in more depth, including security logs. As a brief introduction, security logs will record who took an action, what action they took, where they took it, and when they took it. In other words, if users access a file server over a network, the audit log entries show the user identities, when they accessed the server, what server they accessed, and what computer they used to access the server. Users would not be able to refute the recorded action because auditing provides non-repudiation.

Remember this

You can identify when a user logs on to a local system and when a user accesses a remote system by monitoring account logon events. Configuring account logon monitoring is an important security step for

system monitoring.

Credential Management

A credential is a collection of information that provides an identity (such as a username) and proves that identity (such as with a password). Over time, users often have multiple credentials that they need to remember, especially when they access many web sites. Credential management systems help users store these credentials securely. The goal is to simplify credential management for users, while also ensuring that unauthorized personnel do not have access to the users' credentials.

As an example of a credential management system, Windows 7 includes the Credential Manager, accessible from Control Panel. Users are able to add credentials into the Credential Manager, which stores them securely in special folders called vaults. Then when users access web sites needing credentials, the system automatically retrieves the credentials from the vault and submits them to the web site.

Comparing Access Control Models

Access control ensures that only authenticated and authorized entities can access resources. For example, it ensures that only authenticated users who have been granted appropriate permissions can access files on a server. As stressed in Chapter 1, access control is dependent on accurately identifying users and authenticating them. However, once users have been identified and authenticated, it's possible to grant them access using one of several different models.

You're probably familiar with some of these topics, but the terms role-BAC, rule-BAC, DAC, and MAC may be unfamiliar. By understanding a little more of the underlying design principles, you'll understand why some of the rules are important, and you'll be better prepared to ensure that security principles are followed. The models covered in this section are:

- Role-based access control (role-BAC)
- Rule-based access control (rule-BAC)
- Discretionary access control (DAC)
- Mandatory access control (MAC)

You may notice that CompTIA uses the acronym RBAC for both rule-based access control and role-based access control. For clarity, this chapter uses role-BAC or rule-BAC instead of the ambiguous RBAC.

Often, when using any of the models, you'll run across the following terms:

- **Subjects.** Subjects are typically users or groups that access an object. Occasionally, the subject may be a service that is using a service account to access an object.
- **Objects.** Objects are items such as files, folders, shares, and printers that subjects access. For example, users access files and printers. The access control model (role-BAC, rule-BAC, DAC, or MAC) helps determine how a system grants subjects authorization to objects. Or, said another way, the access control model determines how a system grants users access to files and other resources.

Role-Based Access Control

Role-based access control uses roles to manage rights and permissions for users. This is useful for users within a specific department who perform the same job functions. An administrator creates the roles and then assigns specific rights and permissions to the roles (instead of to the users). When an administrator adds a user to a role, the user has all the rights and permissions of that role.

Using Roles Based on Jobs and Functions

Imagine your organization has several departments such as Accounting, Sales, and IT, and each department has a separate server hosting its files. You can create roles of Accounting, Sales, and IT and assign these roles to users based on the department where they work. Next, you'd grant these roles access to the appropriate server. For example, you'd grant the Accounting role to the Accounting server, grant the Sales role to the Sales server, and so on.

Another example of the role-BAC model is Microsoft Project Server. The Project Server can host multiple projects managed by different project managers. It includes the following roles:

- **Administrators.** Administrators have complete access and control over everything on the server, including all of the projects managed on the server.
- **Executives.** Executives can access data from any project held on the server, but do not have access to modify system settings on the server.
- **Project Managers.** Project managers have full control over their own projects, but do not have any control over projects owned by other project managers.
- **Team Members.** Team members can typically report on work that project managers assign to them, but they have little access outside the scope of their assignments.

Microsoft Project Server includes more roles, but you can see the point with these four. Each of these roles has rights and permissions assigned to it, and to give someone the associated privileges, you'd simply add the user's account to the role.

Documenting Roles with a Matrix

Think about the developers of Microsoft Project Server. They didn't just start creating roles. Instead, they did some planning and identified the roles they envisioned in the application. Next, they identified the privileges each of these roles required. It's common to document role-based permissions with a matrix listing all of the job titles and the privileges for each role, as shown in Table 2.1.

Role	Server Privileges	Project Privileges
Administrators	All	All
Executives	None	All
Project Managers	None	All on assigned projects No access on unassigned projects
Team Members	None	Access for assigned tasks Limited views within scope of their assigned tasks No views outside the scope of their assigned tasks

Table 2.1: Role-BAC matrix for Project Server

Role-BAC is also called hierarchy-based or job-based:

- **Hierarchy-based.** In the Project Server example, you can see how top-level roles, such as the Administrators role, have significantly more permissions than lower-level roles, such as the Team Members role. Roles may mimic the hierarchy of an organization.
- **Job-, task-, or function-based.** The Project Server example also shows how the roles are centered on jobs or functions that users need to perform.

Remember this

A role-BAC model uses roles based on jobs and functions. A matrix is a planning document that matches the roles with the required privileges.

Establishing Access with Group-Based Privileges

Administrators commonly grant access in the role-BAC model using roles, and they often implement roles as groups. They assign rights and permissions (privileges) to groups and then add user accounts to the appropriate group. This type of group-based privileges, where access is based on roles or groups, simplifies user administration.


One implementation of the role-BAC model is the Microsoft built-in security groups, and specially created security groups that administrators create on workstations, servers, and within domains.

The Administrators group is an example of a built-in security group. For example, the Administrators group on a local computer includes all of the rights and permissions on that computer. If you want to grant Marge full and complete control to a computer, you could add Marge's user account to the Administrators group on that computer. Once Marge is a member of the Administrators group, she has all the rights and permissions of the group.

Similarly, you can grant other users the ability to back up and restore data by adding their user accounts to the Backup Operators group. Although the built-in groups are very useful, they don't meet all the requirements in most organizations. For example, if your organization wants to separate backup and restore responsibilities, you can create one group that can only back up data and another group that can only restore data.

In Windows domains, administrators often create groups that correspond to the departments of an organization. For example, imagine that Homer,

Marge, and Bart work in the Sales department and need to access data stored in a shared folder named Sales on a network server. An administrator would simplify administration with the following steps, as shown in Figure 2.3:



1. Create a Sales group and add each of the user accounts to the Sales group.
2. Add the Sales group to the Sales folder.
3. Assign appropriate permissions to the Sales group for the Sales folder.

Figure 2.3: Establishing access with groups as roles

If the company adds new salespeople, the administrator creates accounts for them and places their account into the Sales group. These new salespeople now have access to everything assigned to this group. If any users change jobs within the company and leave the Sales department, the administrator removes them from the Sales group. This automatically prevents them from accessing any resources granted to the Sales group. This example shows how to use a group for the Sales department, but you can apply the same steps to any department or group of users.

Without groups, you would use user-assigned privileges. In other words, you would assign all the specific rights and permissions for every user individually. This may work for one or two users, but quickly becomes unmanageable with more users.

As an example, imagine that people within the Sales department need access to 10 different resources (such as files, folders, and printers) within a network. When the company hires a new salesperson, you'd need to assign permissions to these 10 different resources manually, requiring 10 different administrative tasks. If you assign the permissions to the Sales group, you only need to add the new user to one group and you're done.

Groups provide another security benefit. Imagine that a user is promoted out of the Sales department and now works in Marketing. If you have a Marketing group, you can place this user account into the Marketing group and remove the account from the Sales group. Removing the user from the Sales group instantly removes all the rights and permissions from that group. However, if you're not using groups and assign permissions to users directly, you probably won't remember what resources were assigned to the user as a member of the Sales department. Instead, the user will continue to have access to this sales data, violating the principle of least privilege.

Chapter 8 presents information on reviews of user rights and permissions as part of an auditing process. A routine audit of user rights and permissions will detect when users have more rights and permissions than they need. However, using groups as part of a role-based access control model helps prevent the problem.

Remember this

Group-based privileges reduce the administrative workload of access management. Administrators put user accounts into groups, and assign privileges to the groups. Users within a group automatically inherit the privileges assigned to the group.

Fix This Problem!

A colleague shared an extreme example of how the principle of least privilege was violated at his previous job, instead of using group-based privileges to grant access. He was the lone IT administrator, and no matter how much he asked for help, his boss was never able to get him additional staff. The company grew, and he found he was fielding many complaints because users didn't have the access they needed. He knew he needed to improve the administrative model with groups and roles, but without support from his boss, he wasn't able to implement them. The users complained to his boss, who then put more pressure on him. The boss's direction was simply "Fix this problem! I don't want to hear any more of these complaints."

Ultimately, he put all the users into the Domain Admins group. In a Windows Active Directory domain, the Domain Admins group has full rights and permissions to do anything and everything in a domain. Suddenly, all the users had full privileges. This was the equivalent of lighting the fuse on a time bomb. It would only be a matter of time before users purposely or accidentally caused problems with their newfound permissions. Ironically, his boss was happy because the users stopped complaining. Unfortunately, a couple of months later, one of the users found payroll data files on the network and discovered the salaries of other employees. The files and payroll information quickly spread through the company and caused a significant amount of infighting. At that point, his boss's boss wasn't very happy.

It takes a little time and effort to plan and implement role-based access control with groups. However, it reduces overall administration and helps to implement the principle of least privilege.

...

Rule-Based Access Control

Rule-based access control (rule-BAC) uses rules. The most common example is with rules in routers or firewalls. However, more advanced implementations cause rules to trigger within applications, too.

Routers and firewalls use rules within access control lists (ACLs). These rules define the traffic that the devices allow into the network, such as allowing Hypertext Transfer Protocol (HTTP) traffic for web browsers. These rules are typically static. In other words, administrators create the rules and the rules stay the same unless an administrator changes them again.

However, some rules are dynamic. For example, intrusion detection and prevention systems can detect attacks, and then modify rules to block traffic from an attacker. In this case, the attack triggers a change in the rules.

As another example, it's possible to configure user applications with rules. For example, imagine you want to give Homer additional permissions to a database if Marge is absent. You can configure a database rule to trigger a change to these permissions when the system recognizes that Marge is absent.

Remember this

Rule-based access control is based on a set of approved instructions, such as an access control list. Some rule-BAC systems use rules that trigger in response to an event such as modifying ACLs after detecting an attack, or granting additional permissions to a user in certain situations.

Discretionary Access Control

In the DAC model, every object (such as files and folders) has an owner, and the owner establishes access for the objects. Many operating systems, such as Windows and most Unix-based systems, use the DAC model.

A common example of the DAC model is the NT File System (NTFS) used in Windows. NTFS provides security by allowing users and administrators to restrict access to files and folders with permissions. NTFS is based on the DAC model and the following section explains how it uses the DAC model.

SIDs and DACLs

Microsoft systems identify users with security identifiers (SIDs), though you will rarely see a SID. A SID is a long string of characters that is meaningless to most people and may look like this: S-1-5-21-3991871189-223218. Instead of the system displaying the SID, it looks up the name associated with the SID and displays the name. Similarly, Microsoft systems identify groups with a SID.

Every object (such as a file or folder) includes a discretionary access control list (DACL) that identifies who can access it in a system using the DAC model. The DACL is a list of Access Control Entries (ACEs). Each ACE is composed of a SID and the permission(s) granted to the SID. As an example, a folder named Study Notes might have the following permissions assigned:

- Lisa Full Control
- Bart Read
- Maggie Modify

Each of these entries is an ACE and combined they are a DACL. The Viewing a DACL Lab shows how to view the DACL for a folder. You can access the online exercises for this book at <http://gcgapremium.com/labs/>.

The Owner Establishes Access

If users create a file, they are designated as the owner and have explicit control over the file. As the owner, users can modify the permissions on the object by adding user or group accounts to the DACL and assigning the desired permissions.

The DAC model is significantly more flexible than the MAC model described in the next section. MAC has predefined access privileges, and the administrator is required to make the changes. With DAC, if you want to grant another user access to a file you own, you simply make the change, and that user has access.

Remember this

The DAC model specifies that every object has an owner, and the owner has full, explicit control of the object. Microsoft NTFS uses the DAC model.

Beware of Trojans

An inherent flaw associated with the DAC model is the susceptibility to Trojan horses. Chapter 6 presents malware in much more depth, but for this discussion, you should understand some basics related to Trojan horses.

Trojan horses are executable files. They masquerade as something useful, but they include malware. For example, Bart might decide to download and install a program that a friend raved about. After installation, he decides it's not so great and forgets about it. However, the damage is done.

What really happened? When Bart installed the program, it also installed malware. Moreover, if Bart was logged on with administrative privileges when he installed it, the Trojan is able to run with these administrative privileges.

Many organizations require administrators to have two accounts to mitigate the risks associated with Trojans. Most of the time, administrators log on with a regular user account. If the system is infected with malware, the malware has limited permissions assigned to the regular user account. In contrast, if the system is infected with malware while the administrator is logged on with an administrative account, the malware has the elevated permissions of an administrator.

Mandatory Access Control

The MAC model uses labels (sometimes referred to as sensitivity labels or security labels) to determine access. Security administrators assign labels to both subjects (users) and objects (files or folders). When the labels match, the system can grant a subject access to an object. When the labels don't match, the access model blocks access.

Military units make wide use of this model to protect data. You may have seen movies where they show a folder with a big red and black cover page labeled "Top Secret." The cover page identifies the sensitivity label for the data contained within the folder. Users with a Top Secret label (a Top Secret clearance) and a need to know can access the data within the Top Secret folder.

Need to know is an important concept to understand. Just because individuals have a Top Secret clearance, it doesn't mean they should automatically have access to all Top Secret data. Instead, access is restricted based on a need to know. Need to know is similar in concept to the principle of least privilege, but it only applies to data and data permissions. In contrast, the principle of least privilege applies to permissions and rights.

Security-enhanced Linux (SELinux) is one of the few operating systems using the mandatory access control model. SELinux was specifically created to demonstrate how mandatory access controls can be added to an operating system. In contrast, Windows operating systems use the discretionary access control model.

Labels and Lattice

The MAC model uses different levels of security to classify both the users and the data. These levels are defined in a lattice. The lattice can be a complex relationship between different ordered sets of labels. These labels define the boundaries for the security levels.

Figure 2.4 shows how the MAC model uses a lattice to divide access into separate compartments based on a need to know. The lattice starts by defining different levels of Top Secret, Secret, Confidential, and For Official Use. Each of these labels defines specific security boundaries. Within these levels, the lattice defines specific compartments. For example, the Top Secret level includes compartments labeled Nuclear Power Plant, 007, and Happy Sumo.

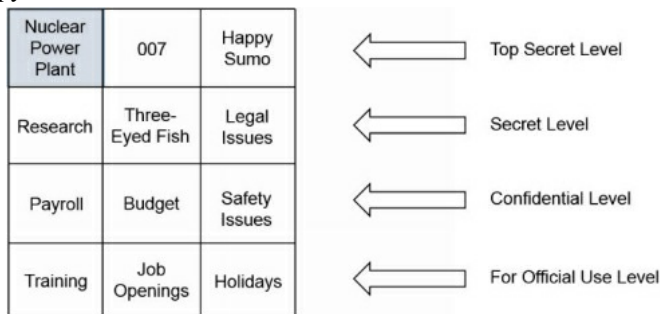


Figure 2.4: MAC model lattice

Imagine that Homer has a Top Secret clearance with a Nuclear Power Plant label. This gives him access to data within the Nuclear Power Plant compartment. However, he does not have access to data in the 007 or Happy Sumo compartment unless he also has those clearances (and associated labels).

Higher-level clearances include lower-level clearances. For example, because Homer has a Top Secret clearance, he can be granted access to Secret and lower-level data. Again though, he will only be able to access data on these lower levels based on his need to know.

As another example, imagine that Lisa has a Secret level clearance. Administrators can grant her access to data on the Secret level and lower levels, based on her need to know. For example, they might grant her access to the Research data by assigning the Research label to her, but not necessarily grant her access to Three-Eyed Fish or Legal Issues data. However, they cannot grant her access to any data on the Top Secret level.

Remember this

The MAC model uses sensitivity labels for users and data. It is commonly used when access needs to be restricted based on a need to know. Sensitivity labels often reflect classification levels of data and clearances granted to individuals.

Establishing Access

An administrator is responsible for establishing access, but only someone at a higher authority can define the access for subjects and objects.

Typically, a security professional identifies the specific access individuals are authorized to access. This person can also upgrade or downgrade the individuals' access, when necessary. Note that the security professional does all this via paperwork and does not assign the rights and permissions on computer systems. Instead, the administrator assigns the rights based on the direction of the security professional.

Multiple approval levels are usually involved in the decision-making process to determine what a user can access. For example, in the military an

officer working in the security professional role would coordinate with higher-level government entities to upgrade or downgrade clearances. These higher-level entities approve or disapprove clearance requests.

Once an individual is formally granted access, a network administrator would be responsible for establishing access based on the clearances identified by the security professional. From the IT administrator’s point of view, all the permissions and access privileges are predefined.

If someone needed different access, the administrator would forward the request to the security professional, who may approve or disapprove the request. On the other hand, the security professional may forward the request to higher entities based on established procedures. This process takes time and results in limited flexibility.

Chapter 2 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Understanding Basic Control Types

- The three primary security control types are technical (implemented with technology), management (using administrative methods), and operational (for day-to-day operations).
- A technical control is one that uses technology to reduce vulnerabilities. Encryption, antivirus software, IDSs, firewalls, and the principle of least privilege are technical controls.
- Management controls are primarily administrative and include items such as risk and vulnerability assessments.
- Operational controls help ensure that day-to-day operations of an organization comply with their overall security plan. Some examples include security awareness and training, configuration management, and change management.
- Preventive controls attempt to prevent security incidents. Examples include system hardening, user training, guards, change management, and account disablement policies.
- Detective controls attempt to detect when a vulnerability has been exploited. Examples include log monitoring, trend analysis, security audits (such as a periodic review of user rights), video surveillance systems, and motion detection systems.
- Corrective controls attempt to reverse the impact of an incident or problem after it has occurred. Examples include active intrusion detection systems, backups, and system recovery plans.
- Deterrent controls attempt to prevent incidents by discouraging threats.
- Compensating controls are alternative controls used when it isn't feasible or possible to use the primary control.

Comparing Physical Security Controls

- Door access control systems should allow personnel to exit without any form of authentication, especially if the systems lose power such as during a fire. Controlled areas such as data centers and server rooms should only have a single entrance and exit point.
- Cipher locks require users to enter a code to open doors. Shoulder surfers can discover the code by watching users enter it, and uneducated users might give out the code to unauthorized personnel. Training reduces these risks.
- A proximity card can electronically unlock a door and helps prevent unauthorized personnel from entering a secure area. By themselves, proximity cards do not identify and authenticate users. Some systems combine proximity cards with PINs for identification and authentication.
- Tailgating occurs when one user follows closely behind another user without using credentials. A mantrap can prevent tailgating. Security guards should be especially vigilant to watch for tailgating in high-traffic areas.
- Security guards are a preventive physical security control and they can prevent unauthorized personnel from entering a secure area. A benefit of guards is that they can recognize people and compare an individual's picture ID for people they don't recognize.
- Closed-circuit television (CCTV) systems provide video surveillance. They provide reliable proof of a person's identity and activity, and can be used to identify individuals entering and exiting secure areas.
- Barricades provide stronger physical security than fences and attempt to deter attackers. Bollards are effective barricades that allow people through, but block vehicles.
- Physical security also includes basic locks. Cable locks secure mobile computers such as laptop computers in a training lab. Server bays include locking cabinets as an additional security measure within a server room. Small devices can be stored in safes or locking office cabinets to prevent the theft of unused resources.

Implementing Logical Access Controls

- The principle of least privilege is a technical control that uses access controls. It specifies that individuals or processes are granted only the rights and permissions needed to perform assigned tasks or functions, but no more.
- Group Policy manages users and computers in a domain, and it is implemented on a domain controller within a domain. Administrators use it to create password policies, lock down the GUI, configure host-based firewalls, and much more.
- Password policies provide a technical means to ensure users employ secure password practices:
 - Password length specifies the minimum number of characters in the password.
 - Password complexity ensures passwords are complex and include at least three of the four character types, such as special characters.
 - Password history remembers past passwords and prevents users from reusing passwords.
 - Minimum password age is used with password history to prevent users from changing their password repeatedly to get back to the original password.

- Maximum password age or password expiration forces users to change their password periodically. When administrators reset user passwords, the password should be immediately expired.
- Password policies should apply to any entity using a password. This includes user accounts and accounts used by services and applications. Applications with internally created passwords should still adhere to the organization's password policy.
- An account disablement policy ensures that inactive accounts are disabled. Accounts for employees who either resign or are terminated should be disabled as soon as possible. Configuring expiration dates on temporary accounts ensures they are disabled automatically.
- Time restrictions can prevent users from logging on or accessing network resources during specific hours.
- Account logon events include when a user logs on locally, and when the user accesses a resource such as a server over the network. These events are logged and can be monitored.

Comparing Access Control Models

- The role-based access control (role-BAC) model uses roles to grant access by placing users into roles based on their assigned jobs, functions, or tasks. A matrix matching job titles with required privileges is useful as a planning document when using role-BAC.
- Group-based privileges are a form of role-BAC. Administrators create groups, add users to the groups, and then assign permissions to the groups. This simplifies administration because administrators do not have to assign permissions to users individually.
- The rule-based access control (rule-BAC) model is based on a set of approved instructions, such as ACL rules in a firewall. Some rule-BAC implementations use rules that trigger in response to an event, such as modifying ACLs after detecting an attack.
- In the discretionary access control (DAC) model, every object has an owner. The owner has explicit access and establishes access for any other user. Microsoft NTFS uses the DAC model, with every object having a discretionary access control list (DACL). The DACL identifies who has access and what access they are granted. A major flaw of the DAC model is its susceptibility to Trojan horses.
- Mandatory access control (MAC) uses security or sensitivity labels to identify objects (what you'll secure) and subjects (users). It is often used when access needs to be restricted based on a need to know. The administrator establishes access based on predefined security labels. These labels are often defined with a lattice to specify the upper and lower security boundaries.

Chapter 2 Practice Questions

1. Which of the following accurately identifies the primary security control classifications?
 - A. Role-based, mandatory, and discretionary
 - B. Technical, management, and operational
 - C. Physical, logical, and technical
 - D. Technical and preventive
2. You need to reduce the attack surface of a web server. Which of the following is a preventive control that will assist with this goal?
 - A. Disabling unnecessary services
 - B. Identifying the initial baseline configuration
 - C. Using hardware locks
 - D. Monitoring logs for trends
3. A security expert is identifying and implementing several different physical deterrent controls to protect an organization's server room. Which of the following choices would BEST meet this objective?
 - A. Using hardware locks
 - B. Utilizing data encryption
 - C. Performing a vulnerability assessment
 - D. Training users
4. You need to secure access to a data center. Which of the following choices provides the BEST physical security to meet this need? (Select THREE.)
 - A. Biometrics
 - B. Cable locks
 - C. CCTV
 - D. Mantrap
5. A security professional needs to identify a physical security control that will identify and authenticate individuals before allowing them to pass, and

restrict passage to only a single person at a time. What should the professional recommend?

- A. Tailgating
- B. Smart cards
- C. Biometrics
- D. Mantrap

6. Your company wants to control access to a restricted area of the building by adding an additional physical security control that includes facial recognition. Which of the following provides the BEST solution?

- A. Bollards
- B. Guards
- C. Palm scanners
- D. Video surveillance

7. Employees access a secure area by entering a cipher code, but this code does not identify individuals. After a recent security incident, management has decided to implement a key card system that will identify individuals who enter and exit this secure area. However, the installation might take six months or longer. Which of the following choices can the organization install immediately to identify individuals who enter or exit the secure area?

- A. Mantrap
- B. Access list
- C. CCTV
- D. Bollards

8. Thieves recently rammed a truck through the entrance of your company’s main building. During the chaos, their partners proceeded to steal a significant amount of IT equipment. Which of the following choices can you use to prevent this from happening again?

- A. Bollards
- B. Guards
- C. CCTV
- D. Mantrap

9. You maintain a training lab with 18 computers. You have enough rights and permissions on these machines so that you can configure them as needed for classes. However, you do not have the rights to add them to your organization’s domain. Which of the following choices BEST describes this example?

- A. Least privilege
- B. Need to know
- C. User-based privileges
- D. Separation of duties

10. Developers in your organization have created an application designed for the sales team. Salespeople can log on to the application using a simple password of 1234. However, this password does not meet the organization’s password policy. What is the BEST response by the security administrator after learning about this?

- A. Nothing. Strong passwords aren’t required in applications.
- B. Modify the security policy to accept this password.
- C. Document this as an exception in the application’s documentation.
- D. Direct the application team manager to ensure the application adheres to the organization’s password policy.

11. You are redesigning your password policy to increase the security of the passwords. Which of the following choices provides the BEST security? (Select TWO.)

- A. Maximum password age
- B. Password complexity
- C. Password history
- D. Password length

12. A company’s account management policy dictates that administrators should disable user accounts instead of deleting them when an employee leaves

the company. What security benefit does this provide?

- A. Ensures that user keys are retained
- B. Ensures that user files are retained
- C. Makes it easier to enable the account if the employee returns
- D. Ensures that users cannot log on remotely

13. You need to create an account for a contractor who will be working at your company for 90 days. Which of the following is the BEST security step to take when creating this account?

- A. Configure history on the account.
- B. Configure a password expiration date on the account.
- C. Configure an expiration date on the account.
- D. Configure complexity.

14. You're asked to identify who is accessing a spreadsheet containing employee salary data. Detailed logging is configured correctly on this file. However, you are unable to identify a specific person who is accessing the file. What is the MOST likely reason?

- A. Shared accounts are not prohibited.
- B. Guest accounts are disabled.
- C. Permissions for the file were assigned to a group.
- D. Account lockout has been enabled.

15. Members of a project team came in on the weekend to complete some work on a key project. However, they found that they were unable to access any of the project data. Which of the following choices is the MOST likely reason why they can't access this data?

- A. Discretionary access control
- B. Time-of-day access control
- C. Rule-based access control
- D. Role-based access control

16. An administrator needs to grant users access to different servers based on their job functions. Which access control model is the BEST choice to use?

- A. Discretionary access control
- B. Mandatory access control
- C. Role-based access control
- D. Rule-based access control

17. Interns from a local college frequently work at your company. Some interns work with the database developers, some interns work with the web application developers, and some interns work with both developers. Interns working with the database developers require specific privileges, and interns working with the web application developers require different privileges. What is the simplest method to meet these requirements?

- A. Use generic accounts.
- B. Create user-based privileges.
- C. Use group-based privileges.
- D. Grant the interns access to the Guest account.

18. Your organization wants to reduce the administrative workload related to account management. Which of the following is the BEST choice?

- A. Implement group-based privileges.
- B. Implement user-based privileges.
- C. Implement the Guest account and Guests group.
- D. Implement periodic reviews of user access.

19. Bart has read access to an accounting database and Lisa has both read and write access to this database. A database application automatically triggers a change in permissions so that Bart has both read and write access when Lisa is absent. What type of access control system is in place?

- A. DAC

- B. MAC
- C. Role-BAC
- D. Rule-BAC

20. Your organization hosts several classified systems in the data center. Management wants to increase security with these systems by implementing two-factor authentication. Management also wants to restrict access to these systems to employees who have a need to know. Which of the following choices should management implement for authorization?

- A. USB token and PIN
- B. Username and password
- C. Mandatory access control
- D. Rule-based access control

Performance-Based Question

Many of the security controls in this chapter can easily be tested in a drag-and-drop or matching type of performance-based question. As long as you know the content, these questions typically aren't any more difficult than a standard multiple-choice question. Here's an example of a performance-based question.

Instructions: You have the following list of controls that you need to use to secure items shown in Figure 2.5:

- Five cable locks
- Four fingerprint readers
- Two proximity badge readers
- One CCTV system
- One mantrap
- One locking cabinet
- One safe

You must use all the items in the list at least once and you must fill all empty boxes in the figure. For example, you must use all five cable locks, not just one cable lock.




Figure 2.5: Matching security controls

Chapter 2 Practice Question Answers

1. **B.** Security controls are classified as technical (implemented by technical means), management (implemented administratively), and operational (for day-to-day operations). Access control methods are role-based, rule-based, mandatory, and discretionary. Physical and logical are not terms used to describe security control classifications, even though some controls are physical and some are logical. Although technical is a security control classification, preventive refers to a security control goal.
2. **A.** Disabling unnecessary services is one of several steps you can take to harden a server and it is a preventive control. Identifying the initial baseline configuration is useful to determine the security posture of the system, but by itself it doesn't prevent attacks. Hardware locks are useful to protect a server room where a web server operates, but it doesn't reduce the attack surface. Monitoring logs and trend analysis are detective controls, not preventive controls.
3. **A.** A hardware lock is a physical security control. It's also a deterrent control because it would deter someone from entering. Data encryption is a technical control designed to protect data and is not a physical security control. A vulnerability assessment is a management control designed to discover vulnerabilities, but it is not a physical control. Training users is an effective preventive control, but it is not a physical control.
4. **A, C, D.** A biometric reader used for access control, a mantrap, and a closed-circuit television (CCTV) system all provide strong physical security for accessing a data center. Cable locks are effective theft deterrents for mobile devices such as laptops, but they don't protect data centers.
5. **D.** A mantrap controls access to a secure area, and only allows a single person to pass at a time. The scenario describes the social engineering tactic

- of tailgating, not the control to prevent it. Some sophisticated mantraps include identification and authorization systems, such as biometric systems or smart cards and PINs. However, biometrics and smart cards used for physical security do not restrict passage to one person at a time unless they are combined with a mantrap.
6. **B.** Security guards can protect access to restricted areas with facial recognition and by checking identities of personnel before letting them in. Bollards are effective barricades to block vehicles, but they do not block personnel. Palm scanners are effective biometric access devices, but they do not use facial recognition. Video surveillance can monitor who goes in and out of an area, but it cannot control the access.
7. **C.** Closed-circuit television (CCTV) or a similar video surveillance system can monitor the entrance and record who enters and exits the area. A mantrap prevents tailgating, but it doesn't necessarily identify individuals. An access list is useful if a guard is identifying users and allowing access based on the access list, but the access list does not identify users. Bollards are a type of barricade that protects building entrances.
8. **A.** Bollards are effective barricades that can block vehicles. Guards can restrict access for personnel, but they cannot stop trucks from ramming through a building. Closed-circuit television (CCTV) or a similar video surveillance system can monitor the entrance, but it won't stop the attack. Mantraps prevent tailgating, but they most likely won't stop a truck.
9. **A.** When following the principle of least privilege, individuals have only enough rights and permissions to perform their job, and this is exactly what is described in this scenario. Need to know typically refers to data and information rather than the privileges required to perform an action, such as adding computers to a domain. User-based privileges refer to giving permissions to individual users rather than groups, and this question doesn't address either user-based privileges or group-based privileges. Separation of duties is a principle that prevents any single person or entity from being able to complete all the functions of a critical or sensitive process, and it isn't addressed in this question either.
10. **D.** The application should be recoded to adhere to the company's password policy, so the best response is to direct the application team manager to do so. Application passwords should be strong and should adhere to an organization's security policy. It is not appropriate to weaken a security policy to match a weakness in an application. Nor is it appropriate to simply document that the application uses a weak password.
11. **B, D.** Password complexity and password length provide the best security. Complexity requires a mix of uppercase and lowercase letters, numbers, and special characters. Length requires a minimum number of characters in the password. Maximum password age requires users to change their password regularly, but by itself allows simple or short passwords. Password history prevents users from reusing passwords.
12. **A.** User accounts typically have security keys associated with them. These keys are retained when the account is disabled, but they are no longer accessible when the account is deleted. By disabling the account, it helps ensure that access to files is retained, but it does not directly retain user files. Employees who leave are not expected to return, so this policy has nothing to do with making it easier to enable an account when they return. Users will not be able to use the accounts locally or remotely if they are disabled or deleted, which is a primary reason to have an account management policy.
13. **C.** When creating temporary accounts, it's best to configure expiration dates so that the system will automatically disable the accounts on the specified date. History, password expiration, and complexity all refer to password policy settings. However, it's rare to configure a specific password policy on a single account.
14. **A.** The most likely reason of those given is that shared accounts are not prohibited, allowing multiple users to access the same file. For example, if the Guest account is enabled and used as a shared account by all users, the logs will indicate the Guest account accessed the file, but it won't identify specific individuals. It doesn't matter how permissions are assigned in order for a log to identify who accessed the file. Account lockout stops someone from guessing a password, but it doesn't affect file access logs.
15. **B.** A time-of-day access control restricts access based on the time of day. It is sometimes used to prevent employees from logging on or accessing resources after normal work hours and during weekends. None of the other options restrict access-based dates or times.
16. **C.** The role-based access control model is the best choice for assigning access based on job functions. A discretionary access control model specifies that every object has an owner and owners have full control over objects, but it isn't related to job functions. Mandatory access control uses labels and a lattice to grant access rather than job functions. A rule-based access control model uses rules that trigger in response to events.
17. **C.** Using group-based privileges is the best choice to meet the needs of this scenario. For example, you can create a DB_Group and a Web_Group, assign appropriate privileges to the groups, and add intern accounts to the groups based on their assignments. User-based privileges take too much time to manage because you'd have to implement them separately. Generic accounts such as the Guest account should not be used.
18. **A.** Group-based privileges reduce the administrative workload related to account management because privileges are assigned to groups that share common responsibilities. User-based privileges are extremely tedious and time consuming because privileges are assigned to all users individually. Generic accounts such as Guest should not be used. Implementing periodic user access reviews is a best practice to ensure accounts are managed properly, but they do not reduce the administrative workload.
19. **D.** A rule-based access control system (rule-BAC) is in place in this scenario with a rule designed to trigger a change in permissions based on an event. The mandatory access control (MAC) model uses labels to identify users and data, and is used in systems requiring a need to know. A discretionary access control (DAC) model does not use triggers. A role-based access control (role-BAC) system uses group-based privileges.
20. **C.** Mandatory access control (MAC) is an access control model that can be used in systems requiring a need to know. It uses labels to identify users and data. If the user has the correct label needed to access the data, the user is authorized access. A USB token and a PIN provide two factors of

authentication, but the question asks what is needed for *authorization*. A username provides identification and a password provides authentication. A rule-based access control system (rule-BAC) uses rules to trigger a change in permissions based on an event, or rules within an access control list (ACL) on hardware devices such as routers.

Performance-Based Question Answer

Figure 2.6 shows the solution to the matching security controls question and the following bullets provide an explanation for the concepts:

- **Unsupervised training lab.** The biggest risk to these laptops is theft and the best theft deterrent for the laptops is cable locks. Because you have five laptops and five cable locks, this uses all of your cable locks for this question.
- **Employee issued laptops.** Of the remaining controls, only fingerprint readers provide protection for the laptops. Additional controls such as encryption and cable locks might also be useful. However, the scenario doesn't include encryption, and it was appropriate to use all the cable locks to protect the laptops in the unsupervised training lab.
- **Server room.** The server room holds much more important data than the lab or the office so it requires the strongest access controls available. In the scenario, you have the following access controls, which you should use with the server room.
 - **Mantrap.** This prevents anyone from tailgating into the server room.
 - **Proximity badge reader.** This provides an added measure of security and you can configure it with the mantrap. (This leaves one proximity badge reader you'll need to use somewhere else.)
 - **CCTV system.** This provides video surveillance and identifies anyone that enters or exits the server room.
 - **Locking cabinets.** These secure the equipment bays (also called equipment cabinets) by locking and preventing access to equipment within them. Without much context, it's difficult to determine if "locking cabinets" refers to equipment cabinets or office cabinets. However, the office has a safe, so an additional locking cabinet isn't necessary.
- **Office.** You can use the leftover controls to secure the office.
 - **Safe.** This is a freebie because the figure shows a picture of a safe.
 - **Proximity badge reader.** You started with two proximity badge readers and used one for the server room and one for the office. You cannot use this to protect employee-issued laptops and it isn't appropriate for an unsupervised training lab that you want people to be able to freely access. The only options are to use one with the server room and one with the office.
 - **Fingerprint reader.** You have one fingerprint reader left after using three with the employee-issued laptops. You can use it on one of the computers here.

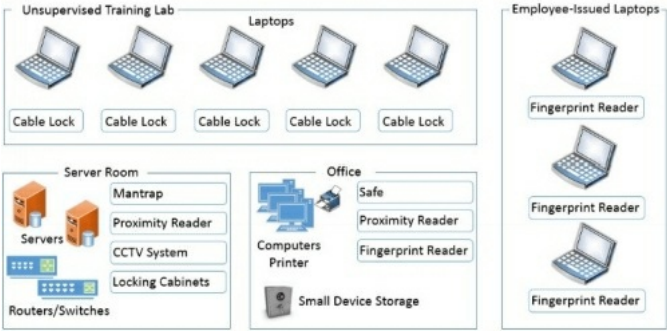


Figure 2.6: Matching security controls solution

Chapter 3

Understanding Basic Network Security

CompTIA Security+ objectives covered in this chapter:

1.1 Implement security configuration parameters on network devices and other technologies.

- Firewalls, Routers, Switches, Proxies, Web security gateways
- UTM security appliances (URL filter, Content inspection, Malware inspection)
- Web application firewall vs. network firewall
- Application aware devices (Firewalls, Proxies)

1.2 Given a scenario, use secure network administration principles.

- Rule-based management, Firewall rules, VLAN management, Secure router configuration, Access control lists, Port Security, 802.1x, Loop protection, Implicit deny, Network separation, Log analysis, Unified Threat Management

1.3 Explain network design elements and components.

- DMZ, Subnetting, VLAN, NAT

1.4 Given a scenario, implement common protocols and services.

- Protocols (IPSec, SNMP, SSH, DNS, TLS, SSL, TCP/IP, FTPS, HTTPS, SCP, ICMP, IPv4, IPv6, FTP, SFTP, TFTP, Telnet, HTTP, NetBIOS)
- Ports (21, 22, 25, 53, 80, 110, 139, 143, 443, 3389)
- OSI relevance

3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques.

- Network security (MAC limiting and filtering, 802.1x, Disabling unused interfaces and unused application service ports)

3.7 Given a scenario, use appropriate tools and techniques to discover security threats and vulnerabilities.

- Tools (Port scanner)

4.3 Given a scenario, select the appropriate solution to establish host security.

- Host-based firewalls

5.2 Given a scenario, select the appropriate authentication, authorization, or access control.

- Authorization (ACLs), Authentication (Implicit deny)

6.2 Given a scenario, use appropriate cryptographic methods.

- Use of algorithms/protocols with transport encryption (SSL, TLS, IPSec, SSH, HTTPS)

**

CompTIA expects prospective CompTIA Security+ exam takers to have at least two years of networking experience. However, even with that amount of experience, there are often gaps in an information technology (IT) professional's or security professional's knowledge. For example, you may have spent a lot of time troubleshooting connectivity but rarely manipulated access control lists (ACLs) on a router or modified firewall rules. This chapter reviews some basic network concepts, including protocols, ports, networking devices, and the Open Systems Interconnection (OSI) reference model. When appropriate, it digs into these topics a little deeper with a focus on security.

Reviewing Basic Networking Concepts

Before you can tackle any of the relevant security issues on a network, you'll need a basic understanding of networking. As a reminder, CompTIA expects you to have a minimum of two years of experience in IT administration; further, CompTIA recommends obtaining the Network+ certification before taking the Security+ exam. Although the Network+ certification isn't required, the knowledge goes a long way in helping you pass the networking portion of the Security+ exam.

This section includes a very brief review of many of the different protocols and networking devices that have a relevance to security. If any of these concepts are completely unfamiliar to you, you may need to pick up a networking book to review them.

Remember this

Networking includes many acronyms, and you'll see a lot of them in this chapter. You can refer to Appendix A, "Acronym List," at the back of the book. It includes a quick reminder of what each acronym represents, along with some key information on the acronym.

Protocols

Networking protocols provide the rules needed for computers to communicate with each other on a network. Some of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocols, such as TCP and IP, provide basic connectivity. Other protocols, such as Hypertext Transfer Protocol (HTTP) and Simple Mail Transfer Protocol (SMTP), support specific types of traffic. This section includes information on common protocols that you'll need to understand for the CompTIA Security+ exam.

It also covers logical ports, including the use of well-known ports. For example, the default port for HTTP is 80. Because ports are so important within IT security, this section lists the associated ports for many of the protocols. Table 3.1 provides a comprehensive listing of all the relevant protocols and ports you should know when taking the CompTIA Security+ exam.

Common TCP/IP Protocols

TCP/IP isn't a single protocol, but a full suite of protocols. Obviously, there isn't room in this book to teach the details of all the TCP/IP protocols. Instead, the purpose of this section is to remind you of some of the commonly used protocols.

If any of these protocols are completely new to you, you might want to do some additional research to ensure you understand the basics. The following sections cover these groups of protocols:

- Basic connectivity protocols
- Encryption protocols
- Application protocols
- Email protocols

Basic Connectivity Protocols

The following list identifies some basic protocols used within the TCP/IP suite for connectivity and testing:

- **TCP.** Transmission Control Protocol (TCP) provides connection-oriented traffic (guaranteed delivery). TCP uses a three-way handshake and Figure 3.1 shows the TCP handshake process. To start a TCP session, the client sends a SYN (synchronize) packet. The server responds with a SYN/ACK (synchronize/acknowledge) packet, and the client completes the third part of the handshake with an ACK packet to establish the connection.




Figure 3.1: TCP handshake process

- **UDP.** User Datagram Protocol (UDP) provides connectionless sessions (without a three-way handshake). ICMP traffic, such as audio and video streaming, uses UDP. Many network-based denial-of-service (DoS) attacks use UDP. TCP/IP traffic is either connection-oriented TCP traffic or connectionless UDP.
- **IP.** The Internet Protocol (IP) identifies hosts in a TCP/IP network and delivers traffic from one host to another using IP addresses. IPv4 uses 32-bit addresses represented in dotted decimal format, such as 192.168.1.100. IPv6 uses 128-bit addresses using hexadecimal code, such as FE80:0000:0000:20D4:3FF7:003F:DE62.
- **ICMP.** Internet Control Message Protocol (ICMP) is used for testing basic connectivity and includes tools such as ping, pathping, and tracert. As an example, ping can check for basic connectivity between two systems. Many DoS attacks use ICMP. Because of how often ICMP is used in attacks, it has become common to block ICMP at firewalls and routers, which disables a ping response. Blocking ICMP prevents attackers from discovering devices in a network with a host enumeration sweep.
- **ARP.** Address Resolution Protocol (ARP) resolves IPv4 addresses to media access control (MAC) addresses. MACs are also called physical addresses, or hardware addresses. TCP/IP uses the IP address to get a packet to a destination network, but once it arrives on the destination network, it uses the MAC address to get it to the correct host. In other words, ARP is required once the packet reaches the destination subnet. ARP poisoning uses ARP packets to give clients false hardware address updates and attackers use it to redirect or interrupt network traffic.
- **NDP.** Neighbor Discovery Protocol (NDP) performs several functions on IPv6. For example, it performs functions similar to IPv4's ARP. It also performs autoconfiguration of device IPv6 addresses and discovers other devices on the network such as the IPv6 address of the default gateway.

Remember this

ARP resolves IPv4 addresses to MAC addresses. NDP performs similar functions on IPv6.

Encryption Protocols

Data in transit is any traffic sent over a network. When data is sent in cleartext, attackers can use a protocol analyzer or sniffer to capture and read it. You can protect the confidentiality of Personally Identifiable Information (PII) and any other sensitive data in transit by encrypting it. Note that you can also encrypt data at rest, which is data stored on any type of medium. Chapter 10, “Understanding Cryptography,” covers several encryption algorithms specifically designed to encrypt data at rest.

The following list identifies several encryption protocols used to encrypt data in transit:

- **SSH.** Secure Shell (SSH) encrypts a wide variety of traffic, such as Secure Copy (SCP) and Secure File Transfer Protocol (SFTP). Unix and Linux administrators often use SSH when remotely administering systems instead of Telnet. SSH can also encrypt TCP Wrappers, a type of access control list used on Linux and Unix systems to filter traffic. When SSH encrypts traffic, it uses TCP port 22.
- **SCP.** Secure Copy (SCP) is based on SSH and is used to copy encrypted files over a network. SCP uses TCP port 22.
- **SSL.** The Secure Sockets Layer (SSL) protocol secures HTTP traffic as Hypertext Transfer Protocol Secure (HTTPS) with the use of certificates. Chapter 10 covers certificates. SSL can also encrypt other types of traffic, such as SMTP and Lightweight Directory Access Protocol (LDAP). SSL uses TCP port 443 when encrypting HTTP, TCP port 465 when encrypting SMTP, and TCP port 636 when encrypting LDAP with SSL.
- **TLS.** The Transport Layer Security (TLS) protocol is the designated replacement for SSL. At this point, you can use TLS instead of SSL in just about any application with the same ports. For example, HTTPS uses TCP port 443 when it uses SSL or TLS. LDAPv2 uses SSL and LDAPv3 uses TLS. Both LDAP with SSL and LDAP with TLS use port 636.
- **IPsec.** Internet Protocol security (IPsec) is used to encrypt IP traffic. It is native to IPv6 but also works with IPv4. IPsec encapsulates and encrypts IP packet payloads and uses Tunnel mode to protect virtual private network (VPN) traffic. IPsec includes two main components: Authentication Header (AH) identified by protocol ID number 51 and Encapsulating Security Payload (ESP) identified by protocol ID number 50. It uses the Internet Key Exchange (IKE) over UDP port 500 to create a security association for the VPN.

Remember this

SSH encrypts a wide variety of traffic and uses port 22 in each implementation. It encrypts File Transfer Protocol (FTP) traffic as SFTP, is used with SCP to copy encrypted files over a network, and can encrypt TCP Wrappers. SSL and TLS encrypt traffic with the use of certificates. You can also use them to encrypt many other types of traffic, such as SMTP and LDAP. IPsec creates secure tunnels for VPNs.

Application Protocols

Application protocols operate on the Application layer of the Open Systems Interconnection (OSI) reference model, or the Application layer of the TCP/IP model. Many of these protocols are responsible for displaying information to the user in a readable format. The following list identifies some of the more commonly used application protocols:

- **HTTP.** Hypertext Transfer Protocol (HTTP) transmits web traffic on the Internet and in intranets. Web servers use HTTP to transmit web pages to clients’ web browsers. Hypertext Markup Language (HTML) is the common language used to display the web pages. HTTP uses TCP port 80.
- **HTTPS.** Hypertext Transfer Protocol Secure (HTTPS) encrypts web traffic to ensure it is secure while in transit. Web browsers commonly indicate that a secure session is using HTTPS by displaying a lock icon and by including HTTPS in the Uniform Resource Locator (URL) field. HTTPS is encrypted with either SSL or TLS and it uses TCP port 443.
- **FTP.** File Transfer Protocol (FTP) uploads and downloads large files to and from an FTP server. By default, FTP transmits data in cleartext, making it easy for an attacker to capture and read FTP data with a sniffer or protocol analyzer. FTP active mode uses TCP port 21 for control signals and TCP port 20 for data. FTP passive mode also uses TCP port 21 for control signals, but it uses a random TCP port for data.
- **SFTP.** Secure File Transfer Protocol (SFTP) is a secure implementation of FTP. It is an extension of Secure Shell (SSH) using SSH to transmit the files in an encrypted format. SFTP transmits data using TCP port 22.
- **FTPS.** File Transfer Protocol Secure (FTPS) is an extension of FTP and uses SSL or TLS to encrypt FTP traffic. Some implementations of FTPS use TCP ports 989 and 990. Notice that the difference between SFTP and FTPS is that SFTP uses SSH and FTPS uses SSL or TLS.
- **TFTP.** Trivial File Transfer Protocol (TFTP) uses UDP and is used to transfer smaller amounts of data, such as when communicating with network devices. Many attacks have used TFTP, but it is not an essential protocol on most networks. Because of this, administrators commonly disable it. TFTP uses UDP port 69.

Remember this

HTTP and HTTPS use ports 80 and 443 and transmit data over the Internet in unencrypted and encrypted formats, respectively. FTP supports uploading and downloading large files to and from an FTP server. FTP uses TCP ports 20 and 21 and TFTP uses UDP port 69. SFTP uses SSH to encrypt FTP traffic and transmits

it using port 22. FTPS uses SSL to encrypt FTP traffic.

- **Telnet.** Telnet is a legacy protocol used to connect to remote systems or network devices over a network. Telnet has a command-line interface, and some administrators use Telnet to connect to routers and make configuration changes. Telnet transmits data in cleartext, making it vulnerable to sniffing attacks. SSH is a more secure alternative than Telnet and applications such as PuTTY have similar functionality but include SSH. PuTTY is the name of a free terminal emulator that many administrators use to connect to remote devices or systems. Telnet uses TCP port 23.
- **SNMP.** Simple Network Management Protocol (SNMP) monitors and manages network devices, such as routers or switches. This includes using SNMP to modify the configuration of the devices or have network devices report status back to a central network management system. SNMP agents installed on devices send information to an SNMP manager via notifications known as traps (sometimes called device traps). The first version of SNMP had vulnerabilities, such as passing passwords across the network in cleartext. SNMPv2 and SNMPv3 are much more secure. SNMP uses UDP port 161. SNMP sends traps (error messages and notifications) on UDP port 162.
- **NetBIOS.** Network Basic Input/Output System (NetBIOS) is a name resolution service for NetBIOS names on internal networks. NetBIOS also includes session services for both TCP and UDP communication. NetBIOS uses UDP ports 137 and 138, and TCP port 139. It can use TCP port 137, but rarely does.
- **LDAP.** Lightweight Directory Access Protocol (LDAP) is the language used to communicate with directories such as Microsoft Active Directory or Novell Netware Directory Services (NDS). LDAP provides a single location for object management and it uses TCP port 389. LDAP can be encrypted with either TLS or SSL and uses port 636 when encrypted.
- **Kerberos.** Kerberos is the authentication protocol used in Windows domains and some Unix environments. It uses a Key Distribution Center (KDC) to issue timestamped tickets. Kerberos uses UDP port 88. Chapter 1, “Mastering Security Basics,” covered Kerberos in more depth.
- **Microsoft SQL Server.** SQL Server is a server application that hosts databases accessible from web servers and a wide array of applications. SQL Server uses port 1433 by default.
- **RDP.** Administrators and clients use Remote Desktop Protocol (RDP) to connect to other systems from remote locations. Microsoft uses RDP in different services such as Remote Desktop Services and Remote Assistance. RDP uses either port TCP 3389 or UDP 3389.

Remember this

Telnet is a legacy protocol administrators have used to connect to remote systems. It uses TCP port 23 and sends data in cleartext. SSH is a more secure alternative than Telnet. Administrators use SNMP to manage and monitor network devices and SNMP uses UDP ports 161 and 162. NetBIOS is used on internal networks and uses ports 137–139. Microsoft SQL Server hosts databases and uses port 1433. Kerberos uses UDP port 88. RDP is used to connect to remote systems and uses port 3389.

Application Protocols

Some common protocols used for email include:

- **SMTP.** Simple Mail Transfer Protocol (SMTP) transfers email between clients and SMTP servers. SMTP uses TCP port 25. SMTP with SSL or TLS uses TCP port 465.
- **POP3.** Post Office Protocol v3 (POP3) transfers emails from servers down to clients. POP3 uses TCP port 110. POP3 with SSL or TLS uses TCP port 995.
- **IMAP4.** Internet Message Access Protocol version 4 (IMAP4) is used to store email on an email server. IMAP4 allows a user to organize and manage email in folders on the server. IMAP4 uses TCP port 143. IMAP4 with SSL or TLS uses TCP port 993.

Remember this

SMTP sends email on TCP port 25, POP3 receives email on port 110, and IMAP4 uses port 143.

IPv4

IPv4 uses 32-bit IP addresses expressed in dotted decimal format. For example, the IPv4 IP address of 192.168.1.5 is four decimals separated by periods or dots. You can also express the address in binary form with 32 bits.

All Internet IP addresses are public IP addresses, and internal networks use private IP addresses. Public IP addresses are tightly controlled. You can't just use any public IP address. Instead, you must either purchase or rent it. Internet Service Providers (ISPs) purchase entire ranges of IP addresses and issue them to customers. If you access the Internet from home, you are very likely receiving a public IP address from an ISP.

Routers on the Internet include rules to drop any traffic that is coming from or going to a private IP address, so you cannot use private IP addresses on the Internet. RFC 1918 specifies the following private address ranges:

- **10.x.y.z.** 10.0.0.0 through 10.255.255.255
- **172.16.y.z–172.31.y.z.** 172.16.0.0 through 172.31.255.255

- **192.168.y.z.** 192.168.0.0 through 192.168.255.255

Subnetting

Subnetting divides a single range of classful IP addresses into two or more smaller ranges of IP addresses. Administrators do this to isolate traffic and increase efficiency. You don't need to know how to subnet for the CompTIA Security+ exam, but you should be familiar with the concept and how administrators use subnetting to isolate systems on different subnets. Additionally, you should be able to identify valid IP addresses for computers within a subnet.

In case you don't remember this from your networking studies, the three primary IP classes are:

- **Class A.** 0.0.0.0 through 127.255.255.255
- **Class B.** 128.0.0.0 through 191.255.255.255
- **Class C.** 192.0.0.0 through 223.255.255.255

Imagine you have multiple users on a single Class C network. Some of the users may be running applications that stream audio and video across the network. A second group of users may regularly upload and download data via the Internet. A third group may upload and download files back and forth to servers on the network, and a fourth group could be users with just occasional access to the network. You can subnet this Class C network into four smaller subnets and isolate the traffic between the groups.

Imagine that the original Class C network is 192.168.1.0 with a subnet mask of 255.255.255.0. It could hold 254 host addresses (192.168.1.1 through 192.168.1.254). You can subnet this into four smaller subnets with each one using a subnet mask of 255.255.255.192 as follows:

- **Subnet 1.** 192.168.1.1 through 192.168.1.62—use for streaming audio and video.
- **Subnet 2.** 192.168.1.65 through 192.168.1.126—use for upload and download of files on the Internet.
- **Subnet 3.** 192.168.1.129 through 192.168.1.190—use for upload and download of files to internal servers.
- **Subnet 4.** 192.168.1.193 through 192.168.1.254—use for regular users.

After dividing the network into the four subnets, you increase the efficiency by reducing collisions on each individual network. This effectively improves the performance of each subnet.

CIDR Notation

Instead of writing out the subnet mask completely, you'll often see it written using Classless Interdomain Routing (CIDR) notation. CIDR notation uses a forward slash (/) followed by a number identifying the number of 1s in the subnet mask. Imagine a subnet mask of 255.255.255.0. It has twenty-four 1s and if you wrote it out in the 32 binary bits, it would be:

- 1111 1111 . 1111 1111. 1111 1111. 0000 0000

With this in mind, you could write an IP address with the same subnet mask in either of these two ways:

- 192.168.100.1, 255.255.255.0
- 192.168.100.1 /24

The IP address ranges in the earlier "Subnetting" section used a subnet mask of 255.255.255.192. This has twenty-six 1s in the subnet mask, so those subnets could be written as:

- 192.168.1.1 /26 through 192.168.1.62 /26
- 192.168.1.65 /26 through 192.168.1.126 /26
- 192.168.1.129 /26 through 192.168.1.190 /26
- 192.168.1.193 /26 through 192.168.1.254 /26

Remember this

Subnetting allows you to divide a classful network into two or more smaller networks. CIDR notation uses a forward slash and a number to identify the subnet mask.

IPv6

Although the number of IP addresses at first seemed inexhaustible, the Internet Assigned Numbers Authority (IANA) assigned the last block of IPv4 addresses in February 2011. To prepare, the Internet Engineering Task Force (IETF) created IPv6, which provides a significantly larger address space than IPv4.

IPv6 supports a significantly larger address space than IPv4 with over 340 undecillion IP addresses. For context, the order is billion, trillion, quadrillion, quintillion, sextillion, septillion, octillion, nonillion, decillion, and undecillion. Everyone will have enough addresses to assign IP addresses to their computers, TVs, mobile phones, refrigerators, coffeemakers, toasters, and anything else they might want to control remotely.

IPv6 uses 128-bit IP addresses expressed in hexadecimal format. For example, the IPv6 IP address of fe80:0000:0000:0000:02d4:3ff7:003f:de62 includes eight groups of four hexadecimal characters, separated by colons. Each hexadecimal character is composed of four bits.

You can simplify IPv6 addresses by omitting leading zeroes in any group of hexadecimal characters, and with zero compression. The IPv6 address looks like this with these two rules:

- Omit leading zeroes: fe80:0:0:2d4:3ff7:3f:de62
- Zero compression: fe80::02d4:3ff7:003f:de62
- Both rules: fe80::2d4:3ff7:3f:de62

The first example omits leading zeroes in five of the groups. Because you know that each group includes four characters, you know that any group with less than four characters is missing the leading zeroes. For example, :2d4 actually represents :02d4.

Zero compression substitutes a string of zeroes with two colons (::). For example, fe80::02d4:3ff7:003f:de62 is the same as fe80:0000:0000:0000:02d4:3ff7:003f:de62. You can only use one double colon in an IPv6 address.

Another benefit of IPv6 over IPv4 is that it has more security built in. For example, the tunneling services provided by IPsec are built in to IPv6 and work natively with it. This allows you to encrypt just about any data in transit relatively easily, including older legacy protocols such as Telnet.

In contrast, IPsec isn't native to IPv4 and it has some compatibility issues. For example, when IPsec passes through a device using Network Address Translation (NAT), NAT breaks IPsec. Although there are ways to work around the issues, IPv6 doesn't have the same problems.

Remember this

IPv6 has a significantly larger address space than IPv4. IPsec is built in to IPv6 and can encrypt any type of IPv6 traffic.

Understanding DNS

The primary purpose of Domain Name System (DNS) is to resolve host names to IP addresses. Systems are constantly querying DNS, though it is usually transparent to users. As an example, imagine that you want to visit <http://blogs.getcertifiedgetahead.com/>. You enter the URL into your web browser or click a link on a page and your system queries a DNS server for the site's IP address. Figure 3.2 shows what is occurring between your system and DNS. DNS uses UDP port 53 for these types of queries.




Figure 3.2: A basic DNS query

DNS servers host data in zones, which you can think of as a database. Zones include multiple records, including the following:

- **A.** Also called a host record. This record holds the host name and IPv4 address and is the most commonly used record in a DNS server. A DNS client queries DNS with the name using a forward lookup request, and DNS responds with the IPv4 address from this record.
- **AAAA.** This record holds the host name and IPv6 address. It's similar to an A record except that it is for IPv6.
- **PTR.** Also called a pointer record. It is the opposite of an A record. Instead of a DNS client querying DNS with the name, the DNS client queries DNS with the IP address. When configured to do so, the DNS server responds with the name. PTR records are optional so these reverse lookups do not always work.
- **MX.** Also called mail exchanger. An MX record identifies a mail server used for email. The MX record is linked to the A record or AAAA record of a mail server.
- **CNAME.** A canonical name, or alias, allows a single system to have multiple names associated with a single IP address. For example, a server named Server1 in the domain GetCertifiedGetAhead.com might have an alias of FileServer1 in the same domain.

Most DNS servers on the Internet run Berkeley Internet Name Domain (BIND) software and run on Unix or Linux servers. Internal networks can use BIND, but in Microsoft networks, most DNS servers use the Microsoft DNS software.

Occasionally, DNS servers share information with each other in a process known as a zone transfer. In most cases, a zone transfer only includes a small number of updated records. However, some transfers include all the records in the zone. DNS servers use TCP port 53 for zone transfers. In contrast, name resolution queries use UDP port 53.

It's easy to overlook the amount and value of information available in a DNS server. However, if attackers are able to access zone data on an internal DNS server, they can map out an entire network. It lists all the names and IP addresses, and identifies some specific servers such as mail servers. Because of this, DNS administrators configure DNS servers to use secure zone transfers to prevent unauthorized modifications of zone data, and to prevent unauthorized zone transfers.

Remember this

DNS zones include records such as A records for IPv4 addresses and AAAA records for IPv6 addresses. DNS uses TCP port 53 for zone transfers and UDP port 53 for DNS client queries. Most Internet-based DNS servers run BIND software on Unix or Linux servers, and it's common to configure DNS servers to only use secure zone transfers.

Understanding and Identifying Ports

Ports are logical numbers used by TCP/IP to identify what service or application should handle data received by a system. Both TCP and UDP use ports with a total of 65,536 TCP ports (0 to 65,535) and 65,536 UDP ports (0 to 65,535). Administrators open ports on firewalls and routers to allow the associated protocol into or out of a network. For example, HTTP uses port 80, and an administrator allows HTTP traffic by opening port 80.

The Internet Assigned Numbers Authority (IANA) maintains a list of official port assignments that you can view at <http://www.iana.org/assignments/port-numbers>. IANA divided the ports into three ranges, as follows:

- **Well-known ports: 0–1023.** IANA assigns port numbers to commonly used protocols in the well-known ports range.
- **Registered ports: 1024–49,151.** IANA registers these ports for companies as a convenience to the IT community. A single company may register a port for a proprietary use, or multiple companies may use the same port for a specific standard. As an example, Microsoft SQL Server uses port 1433 for database servers, Layer 2 Tunneling Protocol (L2TP) uses port 1701, and Point-to-Point Tunneling Protocol (PPTP) uses port 1723.
- **Dynamic and private ports: 49,152–65,535.** These ports are available for use by any application. Applications commonly use these ports to temporarily map an application to a port. These are also called ephemeral ports, indicating that they are short lived.

Although virtually all of the ports are subject to attack, most port attacks are against the well-known ports. Port scanners often simply check to see if a well-known port is open. For example, SMTP uses the well-known port 25, so if port 25 is open, the system is likely running SMTP.

IT personnel who regularly work with routers and firewalls can readily tell you which protocol is associated with which well-known port, such as 20, 21, 22, 23, 25, 80, or 443. The reason is that they use these ports to allow or block traffic.

For example, an administrator can close port 23 to block all Telnet traffic into a network. The router then ignores traffic on port 23 instead of forwarding it. Similarly, an administrator can close port 1433 to block database traffic to a Microsoft SQL server. On the other hand, the administrator can open port 25 to allow SMTP traffic.

Although ports are second nature to router and firewall administrators, they might not be so familiar to you. If you don't work with the ports often, you'll need to spend some extra time studying to ensure you're ready for the exam.

Combining the IP Address and the Port

At any moment, a computer could be receiving dozens of packets. Each of these packets includes a destination IP address and a destination port. TCP/IP uses the IP address to get the packet to the computer. The computer then uses the port number to get the packet to the correct service, protocol, or application that can process it.

For example, if the packet has a destination port of 80 (the well-known port for HTTP), the system passes the packet to the process handling HTTP. It wouldn't do much good to pass an SMTP email packet to the HTTP service or send an HTTP request packet to the SMTP service.

IP Address Used to Locate Hosts

Imagine that the IP address of *GetCertifiedGetAhead.com* is 72.52.206.134, and the address assigned to your computer from your ISP is 70.150.56.80. TCP/IP uses these IP addresses to get the packets from your computer to the web server and the web server's answer back to your computer.

There's a lot more that occurs under the hood with TCP/IP (such as DNS, NAT, and ARP), but the main point is that the server's IP address is used to get the requesting packet from your computer to the server. The server gets the response packets back to your computer using your IP address (or the IP address of your NAT server).

Server Ports

Different protocols are enabled and running on a server. These protocols have well-known or registered port numbers, such as port 22 for SSH, 23 for Telnet, 80 for HTTP, 443 for HTTPS, and so on. When the system receives traffic with a destination of port 80, the system knows to send it to the service handling HTTP.

Any web browser knows that the well-known port for HTTP is 80. Even though you don't see port 80 in the URL, it is implied as <http://GetCertifiedGetAhead.com:80>. If you omit the port number, HTTP uses the well-known port number of 80 by default.

Popular web servers on the Internet include Apache and Internet Information Services (IIS). Apache is free and runs on Unix or Linux systems. Apache can also run on other platforms, such as Microsoft systems. IIS is included in Microsoft Server products, such as Windows Server 2008 and Windows Server 2012. All of these web servers use port 80 for HTTP. When the server receives a packet with a destination port of 80, the server sends the packet to the web server application (Apache or IIS) that processes it and sends back a response.

Client Ports

TCP/IP works with the client operating system to maintain a table of client-side ports. This table associates port numbers with different applications that are expecting return traffic. Client-side ports start at port 49,152 and increment up to 65,535. If the system uses all the ports between 49,152 and 65,535 before being rebooted, it'll start over at 49,152.

When you use your web browser to request a page from a site, your system will record an unused client port number such as 49,152 in an internal table to handle the return traffic. When the web server returns the web page, it includes the client port as a destination port. When the client receives web page packets with a destination port of 49,152, it sends these packets to the web browser application. The browser processes the packets and displays the page.

Putting It All Together

The previous section described the different pieces, but it's useful to put this together into a single description. Imagine that you decide to visit the web site <http://GetCertifiedGetAhead.com> using your web browser so you type the URL into the browser, and the web page appears. Here are the details of what is happening. Figure 3.3 provides an overview of how this will look and the following text explains the process.

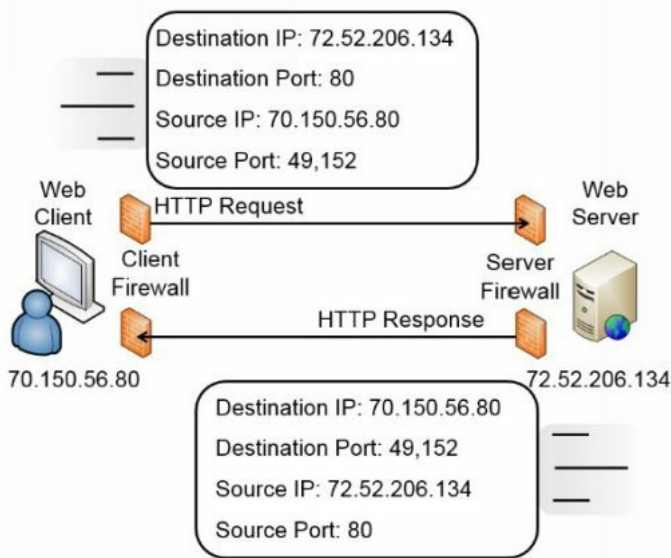


Figure 3.3: Using source and destination ports

Your computer creates a packet with source and destination IP addresses and source and destination ports. It queries a DNS server for the IP address of *GetCertifiedGetAhead.com* and learns that the IP address is 72.52.206.134. Additionally, your computer will use its IP address as the source IP address. For this example, imagine your computer's IP address is 70.150.56.80.

Because the web server is serving web pages using HTTP and the well-known port is used, the destination port is 80. Your computer will identify an unused port in the dynamic and private ports range (a port number between 49,152 and 65,535) and map that port to the web browser. For this example, imagine it assigns 49,152 to the web browser. It uses this as the source port.

At this point, the packet has both destination and source data as follows:

- Destination IP address: 72.52.206.134 (the web server)
- Destination port: 80
- Source IP address: 70.150.56.80 (your computer)
- Source port: 49,152

TCP/IP uses the IP address (72.52.206.134) to get the packet to the *GetCertifiedGetAhead* web server. When it reaches the web server, the server looks at the destination port (80) and determines that the packet needs to go to the web server program servicing HTTP. The web server creates the page and puts the data into one or more return packets. At this point, the source and destinations are swapped because the packet is coming from the server back to you:

- Destination IP address: 70.150.56.80 (your computer)
- Destination port: 49,152
- Source IP address: 72.52.206.134 (the web server)
- Source port: 80

Again, TCP/IP uses the IP address to get the packets to the destination, which is your computer at this point. Once the packets reach your system, it sees that port 49,152 is the destination port. Because your system mapped this port to your web browser, it sends the packets to the web browser, which displays the web page.

The Importance of Ports in Security

Routers, and the routing component of firewalls, filter packets based on IP addresses, ports, and some protocols such as ICMP or IPsec. Because many protocols use well-known ports, you can control protocol traffic by allowing or blocking traffic based on the port.

In the previous example, the client firewall must allow outgoing traffic on port 80. Firewalls automatically determine the client ports used for return traffic, and if they allow the outgoing traffic, they allow the return traffic. In other words, because the firewall allows the packet to the web server on port 80, it also allows the web page returning on the dynamic port of 49,152.

Note that the client firewall doesn't need to allow incoming traffic on port 80 for this to work. The web client isn't hosting a web server with HTTP, so the client firewall would block incoming traffic on port 80. However, the firewall that is filtering traffic to the web server needs to allow incoming traffic on port 80.

You can apply this same principle for any protocol and port. For example, if you want to allow SMTP traffic, you create a rule on the firewall to allow traffic on port 25. Similarly, if you want to block Telnet traffic, you ensure that the firewall blocks port 23.

IT professionals modifying ACLs on routers and firewalls commonly refer to this as opening a port to allow traffic or closing a port to block traffic.

Comparing Ports and Protocol IDs

Ports and protocol identifiers (protocol IDs) are not the same thing, though they are often confused. Well-known ports identify many services or protocols, as discussed previously.

However, many protocols aren’t identified by the port, but instead by the protocol ID. For example, within IPsec, protocol ID 50 indicates the packet is an Encapsulating Security Payload (ESP) packet, and protocol ID 51 indicates it’s an Authentication Header (AH) packet. Similarly, ICMP has a protocol ID of 1, TCP is 6, and UDP is 17.

You can use a protocol ID to block or allow traffic on routers and firewalls just as you can block or allow traffic based on the port. Note that it isn’t accurate to say that you can allow IPsec ESP traffic by opening *port* 50. IANA lists port 50 as a Remote Mail Checking Protocol. However, you can allow IPsec traffic by allowing traffic using protocol ID 50.

Protocol analyzers can capture and examine IP headers to determine the protocol ID and the port, as well as read any unencrypted data. Chapter 8, “Managing Risk,” covers protocol analyzers in more depth.

...

Memorize These Ports

There are 1,024 well-known ports, but you don’t need to know them all. However, at a minimum, you should know the ports listed in Table 3.1. It lists all of the ports specifically mentioned in the CompTIA objectives along with several others you should know when preparing for the exam. Some of these ports are outside the well-known ports range, and some of the protocols are discussed in other chapters in this book. However, you can use this table as a single source to memorize all of the relevant ports.

Protocol	Port	Protocol	Port
FTP data port (active mode)	TCP 20	NetBIOS (TCP rarely used)	TCP/UDP 137
FTP control port	TCP 21	NetBIOS	UDP 138
SSH	TCP 22	NetBIOS	TCP 139
SCP (uses SSH)	TCP 22	IMAP4	TCP 143
SFTP (uses SSH)	TCP 22	LDAP	TCP 389
Telnet	TCP 23	HTTPS	TCP 443
SMTP	TCP 25	SMTP SSL/TLS	TCP 465
TACACS+	TCP 49	IPsec (for VPN with IKE)	UDP 500
DNS name queries	UDP 53	LDAP/SSL	TCP 636
DNS zone transfers	TCP 53	LDAP/TLS	TCP 636
TFTP	UDP 69	IMAP SSL/TLS	TCP 993
HTTP	TCP 80	POP SSL/TLS	TCP 995
Kerberos	UDP 88	L2TP	UDP 1701
POP3	TCP 110	PPTP	TCP 1723
SNMP	UDP 161	Remote Desktop Protocol (RDP)	TCP/UDP 3389
SNMP trap	UDP 162	Microsoft SQL Server	TCP 1433

Table 3.1: Some commonly used well-known ports

Some topics just require you to memorize information, and port numbers is one of those topics. It’s worth your time to memorize these port numbers and their related protocol. Don’t be surprised if you see questions such as “What is the default port for _____?” If you have the table memorized, these questions will be trivial to you.

When you take the CompTIA Security+ exam, you can write down your own notes as soon as you start. Many successful test takers memorize the ports in this table and write down the table as their very first action when they start the exam. Later, when they come across a question that requires the knowledge of a port number, it’s as simple as looking at their notes. If you don’t know these ports now, practice writing this table from memory so you’re ready when it’s time for the live exam.

Remember this

Administrators use ports to identify traffic they want to allow or block. For example, SSH, SCP, and SFTP use TCP port 22 by default. So, by configuring a firewall to allow traffic on port 22, they are allowing SSH, SCP, and SFTP traffic. Memorize the ports in Table 3.1 so that you can answer CompTIA Security+ port-related questions very easily.

Understanding Basic Network Devices

Networks connect computing devices together so that users can share resources, such as data, printers, and other devices. Any device with an IP address is a host, but you'll often see them referred to as clients or nodes. Network devices such as hubs or switches connect these hosts together within a network. Routers connect multiple networks together to create larger and larger networks.

When discussing the different network devices, it's important to remember two primary methods IPv4 uses when addressing TCP/IP traffic:

- **Unicast.** One-to-one traffic. One host sends traffic to another host, using a destination IP address. Only the host with the destination IP address will process the packet.
- **Broadcast.** One-to-all traffic. One host sends traffic to all other hosts on the subnet, using a broadcast address such as 255.255.255.255. Every host that receives broadcast traffic will process it. Hubs and switches pass broadcast traffic between their ports, but routers do not pass broadcast traffic.

Hub

A hub has multiple physical ports used to provide basic connectivity to multiple computers. Hubs commonly have between 4 and 32 physical ports. In an Ethernet network, the hub would have multiple RJ-45 ports used to connect to network interface cards (NICs) on the host computers using twisted-pair cable. Most hubs are active, meaning they have power and will amplify the output to a set level.

Hubs have zero intelligence. Whatever goes in one port goes out all ports on the hub. This presents a security risk because if an attacker installs a protocol analyzer (sniffer) on any computer connected to the hub, the sniffer will capture all the traffic passing through the hub.

As mentioned previously, any traffic sent across the wire in cleartext is subject to sniffing attacks with a protocol analyzer. One way to protect against this is by encrypting the data. Another way of protecting against sniffing attacks is to replace all hubs with switches to limit the amount of traffic that reaches any computer. Many companies specifically restrict the use of hubs in their networks to reduce the risk of sniffers capturing traffic. The following section describes how switches limit traffic to specific computers.

Comparing Ports and Ports

Note that a physical port used by a network device, such as a switch or a router, is completely different from the logical ports discussed previously. You plug a cable into a physical port. A logical port is a number embedded in a packet and identifies services and protocols.

This is similar to minute (sixty seconds) and minute (tiny), or like the old joke about the meaning of *secure*. The Secretary of Defense directed members of different services to “secure that building.” Navy personnel turned off the lights and locked the doors. The Army occupied the building and ensured no one could enter. The Marines attacked it, captured it, and set up defenses to hold it. The Air Force secured a two-year lease with an option to buy.

...

Switch

A switch has the ability to learn which computers are attached to each of its physical ports. It then uses this knowledge to create internal switched connections when two computers communicate with each other.

Consider Figure 3.4. When the switch turns on, it starts out without any knowledge other than knowing it has four physical ports. Imagine that the first traffic is the beginning of a TCP/IP conversation between Lisa's computer and Homer's computer.




Figure 3.4: Switch

When Lisa's computer sends the first packet, it includes the MAC address of the destination computer. However, because the switch doesn't know which port Homer's computer is connected to, it forwards this first packet to all the ports on the switch.

Included in that first packet is the MAC address of Lisa's computer. The switch logs this information into an internal table. It then directs any future traffic addressed to Lisa's MAC address to port 1, and port 1 only.

When Homer's computer receives the packet, it responds. Embedded in this return packet is the MAC address of Homer's computer. The switch captures Homer's MAC address and logs it with port 4 in the internal table. From here on, any unicast traffic between Lisa's and Homer's computers is internally switched. Switches will internally switch unicast traffic. However, they pass broadcast traffic to all ports.

Security Benefit of a Switch

Most of the previous discussion is basic networking, but what you really need to know is why it's relevant in security. If an attacker installed a sniffer on a computer attached to another port (such as port 3 in Figure 3.4), the sniffer would not capture unicast traffic going through the switch to other ports. If Lisa and Homer are exchanging data on ports 1 and 4, none of the traffic reaches port 3. The sniffer can't capture traffic that doesn't reach the port.

In contrast, if the computers were connected via a hub, unicast traffic goes to all hub ports and the attacker could capture it. This is the main *security* reason why organizations replace hubs with switches. The switch reduces the risk of an attacker capturing data with a sniffer. Of course, switches also increase the efficiency of a network.

Physical Security of a Switch

Many switches have a console port that administrators can use to monitor all traffic. Unlike the normal ports that only see traffic specifically addressed to the port, the monitoring port will see all traffic in or out of the switch. This includes any unicast traffic the switch is internally switching between two regular ports. The monitoring port is useful for legitimate troubleshooting, but if the switch isn't protected with physical security, it can also be useful to an attacker.

Physical security protects a switch by keeping it in a secure area such as in a locked wiring closet. Physical security ensures that attackers don't have physical access to the switch and other network devices.

Loop Protection

In some situations, a network can develop a switching loop or bridge loop problem. The effect is similar to a broadcast storm and it can effectively disable a switch. For example, if a user connects two ports of a switch together with a cable, it creates a switching loop where the switch continuously sends and resends unicast transmissions through the switch. In addition to disabling the switch, it also degrades performance of the overall network.

This is trivial for many network administrators, because most current switches have Spanning Tree Protocol (STP) or the newer Rapid STP (RSTP) installed and enabled. STP and RSTP protect against switching loops. However, if these protocols are disabled, the switch is susceptible to loop problems. The simple solution is to ensure that switches include loop protection such as STP or RSTP.

Spanning Tree Protocol also protects the network against potential attackers. For example, imagine an attacker visits a conference room and has access to RJ-45 wall jacks. If loop protection isn't enabled, he can connect two jacks together with a cable, slowing network performance down to a crawl.

Remember this

Loop protection such as STP or RSTP is necessary to protect against switching loop problems, such as those caused when two ports of a switch are connected together.

VLAN

A virtual local area network (VLAN) uses a switch to group several different computers into a virtual network. You can group the computers together based on departments, job function, or any other administrative need. This provides security because you're able to isolate the traffic between the computers in the VLAN.

Normally, a router would group different computers onto different subnets, based on physical locations. All the computers in a routed segment are located in the same physical location, such as on a specific floor or wing of a building.

However, a single switch can create multiple VLANs to separate the computers based on logical needs rather than physical location. Additionally, administrators can easily reconfigure the switch to add or subtract computers from any VLAN if the need arises.

For example, a group of users who normally work in separate departments may begin work on a project that requires them to be on the same subnet. You can configure a switch to logically group these workers together, even if the computers are physically located on different floors or different wings of the building. When the project is over, you can simply reconfigure the switch to return the network to its original configuration.

Similarly, you can use a single switch with multiple VLANs to separate users. For example, if you want to separate the traffic between the HR department and the IT department, you can use a single switch with two VLANs. The VLANs logically separate all the computers between the two different departments, even if the computers are located close to each other.

Remember this

You can create multiple VLANs with a single switch. A VLAN can logically group several different computers together, or logically separate computers, without regard to their physical location.

Port Security

Port security limits the computers that can connect to ports on a switch. At the most basic level, administrators disable unused ports. For example, individual RJ-45 wall jacks in an office lead to specific physical ports on a switch. If the wall jack is not being used, administrators can disable the physical port on the switch. This prevents someone from plugging in a laptop or other computer into the wall jack and connecting to the network.

MAC address filtering is another example of port security. In a simple implementation, the switch remembers the first one or two MAC addresses that connect to a port. It then blocks access to systems using any other MAC addresses. You can also manually configure each port to accept traffic only from a specific MAC address. This limits each port's connectivity to a specific device using this MAC address. This can be very labor intensive, but it provides a higher level of security.

802.1x Port Security

The 802.1x protocol is a port-based authentication protocol and it provides much stronger port security than simply disabling unused ports or using MAC address filtering. It requires users or devices to authenticate when they connect to a specific access point, or a specific physical port, and can be implemented in both wireless and wired networks. It secures the authentication process prior to a client gaining access to a network, and blocks network access if the client cannot authenticate.

You can implement 802.1x as a Remote Authentication Dial-In User Service (RADIUS) or Diameter server. As discussed in Chapter 1, RADIUS provides centralized authentication, and Diameter is an improvement over RADIUS, supporting additional features such as Extensible Authentication Protocol (EAP). 802.1x also supports EAP and can be implemented to require authentication using multiple methods, including digital certificates.

The 802.1x server will prevent rogue devices from connecting to a network. Consider open RJ-45 wall jacks. Although disabling them is a good port security practice, you can also configure an 802.1x server to require authentication for these ports. If clients cannot authenticate, the 802.1x server blocks or restricts access to the network.

It's possible to combine an 802.1x server with other network elements such as a VLAN. For example, imagine you want to provide visitors with Internet access, but prevent them from accessing internal network resources. You can configure the 802.1x server to grant full access to authorized clients, but redirect unauthorized clients to another area of the network via a VLAN.

Chapter 4, "Securing Your Network," covers wireless networks and wireless standards such as Wi-Fi Protected Access (WPA) and Wi-Fi Protected Access II (WPA2). WPA and WPA2 both support Enterprise mode, which uses an 802.1x server for authentication. In contrast, home wireless networks typically use Personal mode without an extra authentication server.

Remember this

Port security includes disabling unused ports and limiting the number of MAC addresses per port. A more advanced implementation is to restrict each physical port to only a single specific MAC address. An 802.1x

server provides port-based authentication, ensuring that only authorized clients can connect to a network. It prevents rogue devices from connecting.

Router

Routers connect multiple network segments together into a single network and route traffic between the segments. As an example, the Internet is effectively a single network hosting billions of computers. Routers route the traffic from segment to segment.

Because routers don't pass broadcasts, they effectively reduce traffic on any single segment. Segments separated by routers are sometimes referred to as broadcast domains. If a network has too many computers on a single segment, broadcasts can result in excessive collisions and reduce network performance. Moving computers to a different segment separated by a router can significantly improve overall performance. Similarly, subnetting networks creates separate broadcast domains.

Cisco routers are popular, but many other brands exist. Most routers are physical devices, and physical routers are the most efficient. However, it's also possible to add routing software to computers with more than one NIC. For example, Windows Server products (such as Windows Server 2008 and Windows Server 2012) can function as routers by adding additional services.

Routers and ACLs

Access control lists (ACLs) are rules implemented on a router (and on firewalls) to identify what traffic is allowed and what traffic is denied. Rules within the ACLs provide rule-based management for the router and control inbound and outbound traffic.

ACLs on routers provide basic packet filtering. They can filter packets based on IP addresses, ports, and some protocols, such as ICMP or IPsec, based on the protocol identifiers:

- **IP addresses and networks.** You can add a rule in the ACL to block access from any single computer based on the IP address. If you want to block traffic from one subnet to another, you can use a rule to block traffic using the subnet IDs. For example, the Sales department may be in the 192.168.1.0/24 network and the Accounting department may be in the 192.168.5.0/24 network. You can ensure traffic from these two departments stays separate with an ACL on a router.
- **Ports.** You can filter traffic based on logical ports. For example, if you want to block HTTP traffic, you can create a rule to block traffic on port 80. Note that you can choose to block incoming traffic, outgoing traffic, or both. In other words, it's possible to allow outgoing HTTP traffic while blocking incoming HTTP traffic.
- **Protocol identifiers.** Many protocols are identified by their protocol IDs. For example, ICMP uses a protocol ID of 1 and many DoS attacks use ICMP. You can block all ICMP traffic (and the attacks that use it) by blocking traffic using this protocol ID. Many automated intrusion protection systems (IPSs) dynamically block ICMP traffic in response to attacks. Similarly, you can restrict traffic to only packets encrypted with IPsec ESP using a rule that allows traffic using protocol ID 50, but blocks all other traffic. PPTP uses protocol ID 47.

Implicit Deny

Implicit deny is an important concept to understand, especially in the context of ACLs. It indicates that all traffic that isn't explicitly allowed is implicitly denied. For example, imagine you configure a router to allow Hypertext Transfer Protocol (HTTP) to a web server. The router now has an explicit rule defined to allow this traffic to the server. If you don't define any other rules, then the implicit deny rule blocks all other traffic.

The implicit deny rule is the last rule in an ACL. Some devices automatically apply the implicit deny rule as the last rule. However, some devices require an administrator to place the rule at the end of the ACL manually. Syntax of an implicit deny rule varies on different systems, but it might be something like Deny Any Any, or Deny All All, where Any indicates any type of traffic and All indicates all traffic.

Remember this

Routers and packet-filtering firewalls perform basic filtering with an access control list (ACL). ACLs identify what traffic is allowed and what traffic is blocked. An ACL can control traffic based on networks, subnets, IP addresses, ports, and some protocols. Implicit deny blocks all access that has not been explicitly granted. Routers and firewalls use implicit deny as the last rule in the access control list.

Firewall

A firewall filters traffic between networks and can filter both incoming and outgoing traffic. In other words, a firewall can ensure only specific types of traffic are allowed into your network and only specific types of traffic are allowed out of your network.

The purpose of a firewall in a network is similar to a firewall in a car. The firewall in a car is located between the engine and passenger compartment. If a fire starts in the engine compartment, the firewall will provide a layer of protection for passengers in the passenger compartment. Similarly, a firewall in a network will try to keep the bad traffic (often in the form of attackers) out of the network.

Of course, an engine has a lot of moving parts that can do damage to people if they accidentally reach into it while it's running. The firewall in a car protects passengers from touching any of those moving parts. Similarly, a network can also block users from going to places that an administrator deems dangerous. For example, uneducated users could inadvertently download damaging files, but many firewalls can block potentially malicious downloads.

Firewalls start with a basic routing capability for packet filtering as described in the “Routers and ACLs” section. More advanced firewalls go beyond simple packet filtering and include advanced content filtering. The “Unified Threat Management” section later in this chapter describes how some security appliances build on basic firewalls by adding in advanced filtering and inspection capabilities.

Host-Based Firewalls

A host-based firewall monitors traffic going in and out of a single host, such as a server or a workstation. It monitors traffic passing through the NIC and can prevent intrusions into the computer via the NIC. Many operating systems include software-based firewalls used as host-based firewalls. For example, Microsoft has included a host-based firewall on operating systems since Windows XP. Additionally, many third-party host-based firewalls are available.

Figure 3.5 shows the host-based Windows Firewall on Windows 7. Notice that you can configure inbound rules to allow or restrict inbound traffic and outbound rules to allow or restrict outbound traffic. The connection security rules provide additional capabilities, such as configuring an IPsec connection in Tunnel or Transport mode to encrypt the traffic.




Figure 3.5: Personal firewall on Windows 7

Linux systems support iptables and many additions such as ipv6tables, arptables, and so on. Generically, administrators commonly refer to these as xtables. You can configure rules within different tables that work similar to how rules within an ACL work.

Personal firewalls provide valuable protection for systems against unwanted intrusions. Many organizations use personal firewalls on each system in addition to network firewalls as part of an overall defense-in-depth strategy.

It's especially important to use personal firewalls when accessing the Internet in a public place. Free Wi-Fi Internet access is often available in public places such as airports, hotels, and many fast-food establishments, such as Starbucks and even McDonald's. However, connecting to a public Wi-Fi hot spot without the personal firewall enabled is risky, and never recommended.

Remember this

Host-based firewalls provide protection for individual hosts such as servers or workstations. A host-based firewall provides intrusion protection for the host. Linux systems support xtables for firewall capabilities. Network-based firewalls are often dedicated servers or appliances and provide protection for the network.

Network-Based Firewalls

A network-based firewall controls traffic going in and out of a network. It does this by filtering traffic based on firewall rules and allows only authorized traffic to pass through it. Most organizations include at least one network-based firewall at the boundary between their internal network and the Internet.

The network-based firewall is usually a dedicated system with additional software installed to monitor, filter, and log traffic. For example, a

popular network-based firewall used in many larger environments is Sidewinder. This is a dedicated server with proprietary firewall software installed. A network-based firewall would have two or more network interface cards (NICs) and all traffic passes through the firewall. Many network-based firewalls are dedicated servers or appliances.

Firewall Rules

Firewalls use rules implemented as ACLs to identify allowed and blocked traffic. This is similar to how a router uses rules. Firewalls use an implicit deny strategy to block all traffic that is not explicitly allowed. Although rules within ACLs look a little different depending on what hardware you’re using, they generally take the following format:

Permission Protocol Source Destination Port

- **Permission.** You’ll typically see this as PERMIT or ALLOW allowing the traffic. Most systems use DENY to block the traffic.
- **Protocol.** Typically, you’ll see TCP or UDP here, especially when blocking specific TCP or UDP ports. If you want to block both TCP and UDP traffic using the same port, you can use IP instead. Using ICMP here blocks ICMP traffic effectively blocking ping and some other diagnostics that use ICMP.
- **Source.** Traffic comes from a source IP address. You can identify a specific IP address to allow or block, or a range of IP addresses. Wildcards such as ANY or ALL include all IP addresses.
- **Destination.** Traffic is addressed to a destination IP address. You can identify a specific IP address to allow or block, or a range of IP addresses, just as you can with the source. Wildcards such as ANY or ALL include all IP addresses.
- **Port or protocol.** Typically, you’ll see the well-known port such as port 80 for HTTP. However, some devices support codes such as www for HTTP traffic. Some systems support the use of keywords such as eq for equal, lt for less than, and gt for greater than. For example, instead of just using port 80, it might indicate eq 80.

An important step when deploying a firewall is to determine what traffic you want to allow. You start by assuming you have an implicit deny rule and then you add exceptions for traffic that you want to allow. You then create a rule for each exception. As an example, imagine you have to create rules on a firewall to meet the following requirements:

1. Allow all HTTP traffic to a web server with an IP of 192.168.1.25.
2. Allow all HTTP and HTTPS traffic to a web server with an IP of 192.168.1.25.
3. Allow DNS queries from any source to a computer with an IP of 192.168.1.10.
4. Block DNS zone transfer traffic from any source to any destination.
5. Block all DNS traffic from any source to any destination.
6. Implement implicit deny.

How many rules would you create? Which protocols would you use? Which ports? The “Firewall Rules Solution” section, at the end of this chapter, shows the solution.

Remember this

Firewalls use a deny any any, deny any, or a drop all statement at the end of the ACL to enforce an implicit deny strategy. The statement forces the firewall to block any traffic that wasn’t previously allowed in the ACL. The implicit deny strategy provides a secure starting point for a firewall.

Web Application Firewall

A *web application firewall (WAF)* is a firewall specifically designed to protect a web application, which is commonly hosted on a web server. In other words, it’s placed between a server hosting a web application and a client. It can be a stand-alone appliance, or software added to another device.

As an example, an organization may host an e-commerce web site to generate revenue. The web server will be placed within a demilitarized zone (DMZ) (discussed later in this chapter), but due to the data that the web server handles, it needs more protection. A successful buffer overflow attack may be able to take the web server down, allow an attacker to access data, or manipulate data.

Chapter 7, “Identifying Advanced Attacks,” covers different types of attacks, including buffer overflow and cross-site scripting attacks. As an example, many buffer overflow attacks start with a series of no operation (NOOP) commands called a NOOP sled or a NOOP ramp. The WAF inspects the contents of traffic to the web server, can detect this malicious content, and blocks it. Similarly, it can detect malicious code in a cross-scripting attack.

Note that you wouldn’t use a WAF in place of a network-based firewall. Instead, it provides an added layer of protection for the web application in addition to the network-based firewall.

Remember this

Web application firewalls provide strong protection for web servers. They protect against several different types of attacks, with a focus on web application attacks such as cross-site scripting attacks.

Advanced Firewalls

Firewall capabilities have advanced significantly over the years and are frequently identified as separate generations. Each new generation includes the capabilities of the previous generation, but adds newer capabilities. The four generations most commonly mentioned are:

- **First generation.** Packet-filtering rules such as those in the previous section were the first generation of firewalls. First-generation firewalls are stateless. In other words, the firewall examines each packet individually and allows or blocks it based on the set of rules in the ACL.
- **Second generation.** Second-generation firewalls added in stateful inspection. In other words, the firewall keeps track of established sessions and inspects traffic based on its state within a session. It blocks traffic that isn't part of an established session.
- **Third generation.** The third generation added application-level firewalls. An application-level firewall is aware of specific commands used in different applications or protocols. For example, a WAF is an application-level firewall that can inspect HTTP traffic and block malicious HTTP traffic.
- **Next generation.** Current network-based firewalls integrate multiple capabilities into a single firewall. As new threats emerge, vendors update the firewalls to adapt. Many firewalls integrated into unified threat management appliances (discussed later in this chapter) are next-generation firewalls.

Firewall Logs and Log Analysis

In addition to filtering traffic, most firewalls have logs you can use to monitor traffic. Firewalls typically allow you to log all allowed traffic, all blocked traffic, or both. A firewall log is often the first place that an administrator might check to investigate a possible intrusion.

Scripts and applications automate the process of reviewing logs. As an example, intrusion detection systems frequently use firewall logs as a source of raw data to help identify intrusions. No matter how the logs are reviewed though, they will include valuable information that might signal an attack, or help an administrator create an audit trail of an attack.

As a simple example, a port scan attack will query different logical ports of an IP address to see what ports are open. Based on what ports are open, the attacker can determine what services or protocols may be running on the server. For example, if port 25 provides a response back to the scanner, it indicates this port is open. Because SMTP uses port 25, this system is very likely running SMTP and may be an email server. The attacker will then send other packets to gather more information. If logging is enabled on the firewall, all of this activity is recorded and it can be used to thwart the attack.

Network Separation

A common network security practice is to use different components to provide network separation. The components you use are dependent on the network infrastructure, but you can do so with routers, VLANs, and firewalls. An extreme practice is to isolate a network completely by ensuring it doesn't have any connectivity with any other network:

- **Routers.** Routers segment traffic between networks using rules within ACLs. Administrators use subnetting to divide larger IP address ranges into smaller ranges.
- **VLANs.** VLANs segment traffic between logical groups of users or computers, regardless of their physical location. You create VLANs with switches.
- **Firewalls.** Firewalls separate network traffic using basic packet-filtering rules and can also use more sophisticated methods to block undesirable traffic.

Protecting the Network Perimeter

Most networks have Internet connectivity, but it's rare to connect a network directly to the Internet. Instead, the perimeter provides a boundary between an internal network (sometimes called an intranet) and the Internet. Boundary protection includes multiple methods to protect the network perimeter.

DMZ

The *demilitarized zone (DMZ)* is a buffered zone between a private network and the Internet. Attackers seek out servers on the Internet, so any server placed directly on the Internet has the highest amount of risk. However, the DMZ provides a layer of protection for these Internet-facing servers.

As an example, Figure 3.6 shows a common network configuration with a DMZ. The DMZ is the area between the two firewalls (FW1 and FW2) and hosts several Internet-facing servers. Many DMZs have two firewalls creating a buffer zone between the Internet and the internal network, as shown in Figure 3.6, though other configurations are possible.

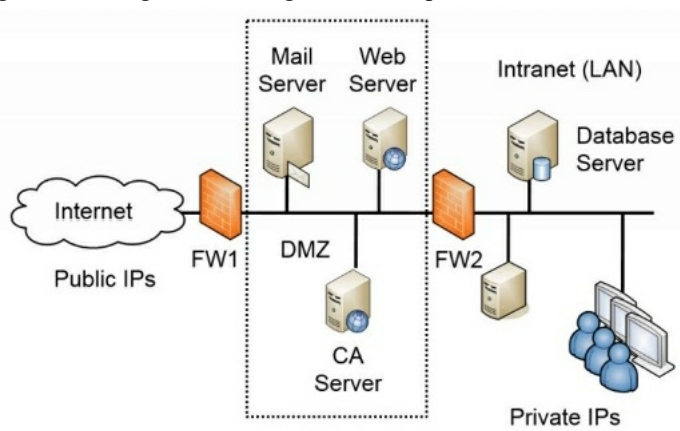


Figure 3.6: Network with DMZ

In this configuration, one firewall separates the DMZ from the Internet. The second firewall separates the DMZ from the internal network. Each firewall includes detailed rules designed to filter traffic and protect both the internal network and the public servers. One way of saying this is that the DMZ provides access to the services hosted in the DMZ, while segmenting access to the internal network.

For example, FW1 can have rules to allow traffic to the servers in the DMZ, but block unsolicited traffic to FW2. The mail server would send and receive email to other email servers on the Internet through port 25 of FW1, and also send and receive email to internal clients through port 25 on FW2. The web server hosts web pages to any Internet users through ports 80 and 443 on FW1, but FW2 blocks these ports. The CA server validates certificates for Internet clients by answering through FW1.

Notice in Figure 3.6 that the intranet includes a database server. The web server may use this to create web pages for an e-commerce site. It could hold product data, customer data, and much more. FW2 allows traffic between the web server (and only the web server) and the database server on port 1433. FW2 would block all other Internet traffic to the database server.

The DMZ can host any Internet-facing server, not just those shown in the figure. Other examples include FTP servers used for uploading and downloading files and virtual private network (VPN) servers used for providing remote access.

Remember this

A DMZ is a buffer zone between the Internet and an internal network. It allows access to services while segmenting access to the internal network. In other words, Internet clients can access the services hosted on servers in the DMZ, but the DMZ provides a layer of protection for the internal network.

Understanding NAT and PAT

Network Address Translation (NAT) is a protocol that translates public IP addresses to private IP addresses and private addresses back to public. You'll often see NAT enabled on an Internet-facing firewall. A commonly used form of NAT is network address and port translation, commonly called Port Address Translation (PAT).

If you run a network at your home (such as a wireless network), the router that connects to the Internet is very likely running NAT. Some of the benefits of NAT include:

- **Public IP addresses don't need to be purchased for all clients.** A home or company network can include multiple computers that can access the Internet through one router running NAT. Larger companies requiring more bandwidth may use more than one public IP address.
- **NAT hides internal computers from the Internet.** Computers with private IP addresses are isolated and hidden from the Internet. NAT provides a layer of protection to these private computers because they aren't as easy to attack and exploit from the Internet.

One of the drawbacks to NAT is that it is not compatible with IPsec. You can use IPsec to create VPN tunnels and use it with L2TP to encrypt VPN traffic. Although there are ways of getting around NAT's incompatibility with IPsec, if your design includes IPsec going through NAT, you'll need to look at it closely.

NAT can be either static NAT or dynamic NAT:

- **Static NAT.** Static NAT uses a single public IP address in a one-to-one mapping. It maps a single private IP address with a single public IP address.
- **Dynamic NAT.** Dynamic NAT uses multiple public IP addresses in a one-to-many mapping. Dynamic NAT decides which public IP address to use based on load. For example, if several users are using the connection on one public IP address, NAT maps the next request to a less-used public IP address.

Remember this

NAT translates public IP addresses to private IP addresses, and private IP addresses back to public. A common form of NAT is Port Address Translation. Dynamic NAT (DNAT) uses multiple public IP addresses while PAT uses a single public IP address.

Proxies

Many networks use proxies, or proxy servers, to forward requests for services (such as HTTP or HTTPS) from clients. They can improve performance by caching content and can restrict users' access to inappropriate web sites by filtering content. A proxy server is located on the edge of the network bordering the Internet and the intranet, as shown in Figure 3.7.




Figure 3.7: Proxy server

All internal clients send requests through the proxy server. The proxy accepts the request, retrieves the content from the Internet, and then returns the data to the client. Most proxy servers only act as a proxy for HTTP and HTTPS. However, proxy servers can also proxy other Internet protocols, such as FTP.

Caching Content for Performance

The proxy server increases the performance of Internet requests by caching each result received from the Internet. Any data that is in the proxy server's cache doesn't need to be retrieved from the Internet again to fulfill another client's request. In this context, *cache* simply means "temporary storage." Cache could be a dedicated area of RAM, or, in some situations, it could also be an area on a high-performance disk subsystem.

As an example, if Lisa retrieves a web page from *GetCertifiedGetAhead.com*, the proxy server would store the result in cache. If Homer later requests the same page, the proxy server retrieves the page from cache and sends it to Homer. This reduces the amount of Internet bandwidth used for web browsing because the page doesn't need to be retrieved again.

Using URL Filters to Restrict Access

Proxy servers can also restrict what users can access with the use of URL filters. A proxy server URL filter examines the requested URL and chooses to allow the request or deny the request.

Many third-party companies sell subscription lists for URL filtering. These sites scour the Internet for web sites and categorize the sites based on what companies typically want to block. Categories may include anonymizers, pornography, gambling, web-based email, and warez sites. *Anonymizers* are sites that give the illusion of privacy on the Internet. Employees sometimes try to use anonymizers to bypass proxy servers, but a proxy server usually detects, blocks, and logs these attempts. Web-based email bypasses the security controls on internal email servers so many organizations block them. Warez sites often host pirated software, movies, MP3 files, and hacking tools.

The subscription list can be loaded into the proxy server, and whenever a user attempts to access a site on the URL filter block list, the proxy blocks the request. Often, the proxy server presents users with a warning page when they try to access a restricted page. Many organizations use this page to remind users of a corporate acceptable usage policy, and some provide reminders that the proxy server is monitoring their online activity.

Proxy servers include logs that record each site visited by users. These logs can be helpful to identify frequently visited sites and to monitor user web browsing activities.

Remember this

A proxy server forwards requests for services from a client. It provides caching to improve performance and reduce Internet bandwidth usage. Proxy servers use URL filters to restrict access to certain sites, and can log user activity.

Unified Threat Management

Unified threat management (UTM) is a single solution that combines multiple security controls. The overall goal of UTMs is to provide better security, while also simplifying management requirements. In many cases, a UTM device will reduce the workload of administrators without sacrificing security.

As IT-based threats first began appearing, security experts created various solutions to deal with each of them. When attackers began releasing malware to infect computers, vendors created antivirus software. Attackers started attacking networks, and in response, security experts developed and steadily improved firewalls. When organizations recognized a need to control what sites users can visit, organizations implemented proxies with URL filters.

Although all of these solutions are effective, they are also complex. Administrators often find it challenging to manage each of these solutions separately. Because of this, UTM security appliances have become quite popular.

Web Security Gateway

A web security gateway is a type of UTM appliance and it can protect against multiple threats. This includes threats from malicious software (malware) coming in as an email attachment, malicious code embedded in web browser pages, and spam. They usually include other firewall capabilities, but their real strength is in content filtering.

Many content filters actively monitor data streams by inspecting the packets in search of malicious code or behaviors. For example, users' email may contain malicious attachments. By inspecting all the packets associated with an email and its attachments, a content filter can detect the malicious content and filter it.

Cisco sells the Web Security Appliance (WSA), which includes several features, including threat defense, content inspection, malware protection, and data loss prevention (DLP) capabilities. Most of these capabilities scan transmissions coming into a network. However, DLP scans data going out of a network. For example, a DLP system can scan all outgoing emails looking for confidential or sensitive information. It would block these emails and identify the user sending them. Chapter 5, "Securing Hosts and Data," discusses DLP in more depth.

UTM Security Appliances

UTM security appliances combine the features of multiple security solutions into a single appliance. For example, a UTM security appliance might include a firewall, antivirus protection, anti-spam protection, URL filtering, and content filtering.

In general, a computer appliance is a hardware device designed to provide a specific solution. For example, spam appliances scan all incoming email and strip off spam. The intent of the word *appliance* is to evoke a sense of simplicity. For example, you don't have to know the details of how a toaster works to make toast. Similarly, you don't have to know the details of how a computer appliance operates to use it.

UTM security appliances include multiple capabilities, including:

- **URL filtering.** URL filters within a UTM security appliance perform the same job as a proxy server. They block access to sites based on the URL. It's common to subscribe to a service and select categories to block access to groups of sites. Administrators can also configure URL filters manually to allow or block access to specific web sites. As an example, if an administrator realizes that users are routinely connecting to a peer-to-peer (P2P) file sharing site, the administrator can add the URL to the filter, and block access to that site.
- **Malware inspection.** Malware often comes into a network via spam, or malicious web pages. The malware inspection component of a UTM appliance screens incoming data for known malware and blocks it. Organizations often scan for malware at email servers and at individual systems as part of a layered security or defense-in-depth solution.
- **Content inspection.** Content inspection includes a combination of different content filters similar to a web security appliance. It monitors incoming data streams and attempts to block any malicious content. It can include a spam filter designed to inspect incoming email and reject spam. It can also block specific types of transmissions, such as streaming audio and video, and specific types of files such as Zip files.

Remember this

A web security gateway and a unified threat management appliance both combine multiple security controls into a single appliance. They can inspect data streams and often include URL filtering, malware inspection, and content inspection components.

Web Security Gateway Versus UTM Security Appliance

You might be wondering what the difference is between a UTM security appliance and a web security gateway. I did. The answer is "not much." Web security gateway was the common term for these appliances when they first started appearing. Later, International Data Corporation (IDC) began using the term in some of its market data and analytics reports. As of this writing, you'll see UTM used much more. Even the CompTIA Security+ objectives stress UTMs more than web security gateways.

The key is that they both can include multiple capabilities. The actual capabilities of the appliance are dependent on what the vendor includes. Most vendors create products based on the capabilities that their customers want or need. As threats evolve, needs change, and you'll likely see the capabilities of these appliances increase to address newer threats.

Identifying OSI Relevance

The Open Systems Interconnection (OSI) reference model conceptually divides different networking requirements into seven separate layers. For most people studying for the CompTIA Security+ exam, the OSI model isn't new. However, because it's mostly theoretical, and rarely used in day-to-day maintenance, some of the knowledge often slips away.

The good news is that you don't need to know it as in-depth as you would for other certification exams such as the CompTIA Network+ exam. If you recently studied for Network+, you probably mastered these concepts and this will just be a quick review. Table 3.2 provides a summary of what you need for the CompTIA Security+ exam and the following sections provide additional details.

Layer Number	Layer Name	Devices	Protocols
1	Physical	Cables, hubs	Ethernet, cabling protocols
2	Data Link	Switches	MAC, ARP, NDP, VLANs
3	Network	Router, Layer 3 switch	IPv4, IPv6, IPsec, ICMP
4	Transport		TCP, UDP
5	Session		
6	Presentation		
7	Application	Proxies, application-proxy firewalls, web application firewalls, web security gateways, UTM security appliances	DNS, FTP, FTPS, HTTP, HTTPS, IMAP4, LDAP, POP3, RDP, SCP, SFTP, SMTP, SNMP, SSH, Telnet, and TFTP

Table 3.2: OSI layers, devices, and protocols

Understanding the Layers

As shown in Table 3.2, the OSI model has seven layers. Many people use mnemonics to memorize the layers. For example, “All People Seem To Need Data Processing” works for some people. The first letter in each of the words represents the first letter of the layer. The A in All is for Application, the P in People is for Presentation, and so on. Another common mnemonic is “Please Do Not Throw Sausage Pizza Away” (for Physical, Data Link, Network, Transport, Session, Presentation, and Application).

After mastering the mnemonic, you also need to remember which layer is Layer 1, and which layer is Layer 7. This memory technique may help. You may have heard about a “Layer 8 error.” This is another way of saying “user error” and users interact with applications. In other words, a user on the mythical Layer 8 interacts with applications, which are on Layer 7. I don’t mean to belittle users or user errors—I make my fair share of errors. However, this memory trick has helped me and many other people remember that the Application layer is Layer 7.

The following sections provide a short synopsis of the OSI model. If you’d like to dig deeper, check out the “Open System Interconnection Protocols” section on Cisco’s DocWiki site at http://docwiki.cisco.com/wiki/Open_System_Interconnection_Protocols.

Layer 1: Physical

The Physical layer is associated with the physical hardware. It includes specifications for cable types, such as 1000BaseT, connectors, and hubs. Computing devices such as computers, servers, routers, and switches transmit data onto the transmission medium in a bit stream. This bit stream is formatted according to specifications at higher-level OSI layers.

Layer 2: Data Link

The Data Link layer is responsible for ensuring that data is transmitted to specific devices on the network. It formats the data into frames and adds a header that includes MAC addresses for the source and destination devices. It adds frame check sequence data to the frame to detect errors. This does not support error correction though. The Data Link layer simply discards frames with detected errors. Flow control functions are also available on this layer.

Switches operate on this layer. As a reminder, computer NICs have a MAC assigned and switches map the computer MAC addresses to physical ports on the switch. Systems use ARP to resolve IPv4 addresses to MAC addresses, and NDP to resolve IPv6 addresses to MAC addresses. VLANs are defined on this layer.

Layer 3: Network

The Network layer uses logical addressing in the form of IP addresses at this layer. This includes both IPv4 addresses and IPv6 addresses. Packets identify where the traffic originated (the source IP address) and where it is going (the destination IP address). Other protocols that operate on this layer are IPsec and ICMP. Routers and Layer 3 switches operate on this layer.

Layer 4: Transport

The Transport layer is responsible for transporting data between systems, commonly referred to as end-to-end connections. It provides reliability with error control, flow control, and segmentation of data. TCP and UDP operate on this layer.

Layer 5: Session

The Session layer is responsible for establishing, maintaining, and terminating sessions between systems. In this context, a session refers to an extended connection between two systems sometimes referred to as dialogs or conversations. As an example, if you log on to a web page, the Session layer establishes a connection with the web server and keeps it open while you’re interacting with the web pages. When you close the pages, the Session layer terminates the session.

If you’re like many users, you probably have more than one application open at a time. For example, in addition to having a web browser open, you might have an email application open. Each of these is a different session, and the Session layer manages them separately.

Layer 6: Presentation

The Presentation layer is responsible for formatting the data as needed by the end-user applications. For example, American Standard Code for Information Interchange (ASCII) and Extended Binary Coded Decimal Interchange Code (EBCDIC) are two standards that define codes used to display characters on this layer.

Layer 7: Application

The Application layer is responsible for displaying information to the end user in a readable format. Application layer protocols typically use this layer to determine if sufficient network resources are available for an application to operate on the network.

Note that this layer doesn't refer to end-user applications directly. However, many end-user applications use protocols defined at this layer. For example, a web browser interacts with DNS services to identify the IP address of a web site name. Similarly, HTTP transmits web pages over the Internet on this layer, which are ultimately displayed in a web browser.

Some of the protocols that operate on this layer are DNS, FTP, FTPS, HTTP, HTTPS, IMAP4, LDAP, POP3, RDP, SCP, SFTP, SMTP, SNMP, SSH, Telnet, and TFTP. SCP isn't defined in an RFC so you won't find a definitive source indicating which layer it operates on. However, SCP uses SSH for data transfer and SSH operates on Layer 7. Similarly, RDP is a proprietary protocol and Microsoft doesn't link it to an OSI layer. However, RDP is listed as an Application layer protocol on the TCP/IP model.

Many advanced devices are application aware and operate on all of the layers up to the Application layer. This includes proxies, application-proxy firewalls, web application firewalls, web security gateways, and UTM security appliances.

Firewall Rules Solution

Table 3.3 shows the rules you’d implement to meet the requirements in the “Firewall Rules” section. Additionally, the following list provides explanations for each these requirements. For clarity, the rules are restated here:

- 1. Allow all HTTP traffic to a web server with an IP of 192.168.1.25.
Note that while HTTP traffic typically uses TCP, it can also use UDP. Because of this, IP is used instead of TCP or UDP.
- 2. Allow all HTTP and HTTPS traffic to a web server with an IP of 192.168.1.25.
This requires two rules. One rule allows HTTP traffic by allowing port 80, and the second rule allows HTTPS traffic by allowing port 443.
- 3. Allow DNS queries from any source to a computer with an IP of 192.168.1.10.
DNS name resolution queries use UDP port 53.
- 4. Block DNS zone transfer traffic from any source to any destination.
DNS zone transfers use TCP port 53.
- 5. Block all DNS traffic from any source to any destination.
Using IP blocks both DNS name resolution queries on UDP port 53 and DNS zone transfers on TCP port 53. You could also implement this as two separate rules with one for UDP and one for TCP.
- 6. Implement implicit deny.
The implicit deny rule is always placed last and it blocks any type of traffic from any source to any destination using any port. Note that you could also have omitted rules 4 and 5 and placed the implicit deny rule after rule 3. It would still have met the requirements but wouldn’t have stressed the difference between TCP port 53 and UDP port 53.

Requirement	Permission	Protocol	Source	Destination	Port
1	ALLOW	IP	ANY	192.168.1.25	80
2	ALLOW	IP	ANY	192.168.1.25	80
2	ALLOW	IP	ANY	192.168.1.25	443
3	ALLOW	UDP	ANY	192.168.1.10	53
4	DENY	TCP	ANY	ANY	53
5	DENY	IP	ANY	ANY	53
6	DENY		ANY	ANY	

Table 3.3: Firewall rules

Chapter 3 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Reviewing Basic Networking Concepts

- TCP and UDP default ports identify specific protocols.
- ARP resolves MAC addresses to IPv4 addresses. NDP performs similar functions on IPv6.
- Several encryption protocols encrypt data in transit to protect its confidentiality. They include SSH, FTPS, SFTP, SCP, IPsec, SSL, and TLS.
- SSH uses TCP port 22. SSH encrypts SFTP traffic, SCP traffic, and TCP Wrappers. SSH uses port 22 when encrypting other protocols.
- SSL and TLS can encrypt many protocols, including HTTPS, SMTP, and LDAP. They both utilize certificates. TLS is the designated replacement for SSL.
- HTTP uses port 80 for web traffic. HTTPS encrypts HTTP traffic in transit, and uses port 443.
- FTP can upload and download large files and it uses TCP port 20 for data and TCP port 21 for control signals. It can be secured with SSH (as SFTP) or with SSL (as FTPS).
- Telnet uses TCP port 23. SSH is a more secure alternative than Telnet.
- SNMP is used to monitor and configure network devices and uses notification messages known as traps. SNMP uses UDP ports 161 and 162.
- NetBIOS uses ports 137–139. Kerberos uses UDP port 88.
- Microsoft SQL Server is database software and it uses port 1433.
- RDP is used to remotely connect to systems and it uses port 3389.
- SMTP sends email using TCP port 25. POP3 receives email using TCP port 110. IMAP4 uses TCP port 143.
- IPv6 uses 128-bit addresses and is displayed as eight groups of hexadecimal characters. It provides a significantly larger address space than IPv4 and natively supports IPsec.
- DNS zones include A records for IPv4 addresses and AAAA records for IPv6 addresses. Zone data is updated with zone transfers and secure zone transfers help prevent unauthorized access to zone data. DNS uses TCP port 53 for zone transfers and UDP port 53 for DNS client queries. Most Internet-based DNS servers run BIND software on Linux or Unix servers.
- Table 3.1 identifies key ports and protocols worth memorizing.

Understanding Basic Network Devices

- Switches are used for network connectivity and they map MAC addresses to physical ports.
- Loop protection protects against switching loop problems, such as when a user connects two switch ports together with a cable. Spanning Tree Protocols protect against switching loops.
- VLANs can logically separate computers or logically group computers regardless of their physical location.
- Port security limits access to switch ports. It includes limiting the number of MAC addresses per port and disabling unused ports. You can also manually map each port to a specific MAC address or group of addresses.
- An 802.1x server provides stronger port security with port-based authentication. It prevents rogue devices from connecting to a network, by ensuring that only authorized clients can connect.
- Implicit deny indicates that unless something is explicitly allowed, it is denied. Firewalls often use implicit deny by explicitly allowing some traffic and then implicitly denying all other traffic that is not identified. Anything not explicitly allowed is implicitly denied.
- A host-based firewall helps protect a single system from intrusions. Some Linux systems use iptables or xtables for firewall capabilities.
- A network-based firewall controls traffic going in and out of a network.
- A firewall controls traffic between networks using rules within an ACL. The ACL can block traffic based on ports, IP addresses, subnets, and some protocols.
- Most firewalls use an implicit deny strategy by blocking all traffic that isn't explicitly allowed. This is implemented with a deny all, or deny any rule at the end of the ACL.
- A web application firewall (WAF) protects a web server against web application attacks such as buffer overflow and cross-site scripting attacks.

Protecting the Network Perimeter

- A DMZ provides a layer of protection for servers that are accessible from the Internet.
- NAT translates public IP addresses to private IP addresses, private back to public, and hides IP addresses on the internal network from users on the Internet.
- A proxy server forwards requests for services from a client. It can filter requests based on URLs, cache content, and record users' Internet

activity.

- A unified threat management (UTM) security appliance includes multiple layers of protection, such as URL filters, content inspection, and malware inspection.

Identifying OSI Relevance

- Know the OSI layers by number and name. The lowest layer is the Physical layer and the highest layer is the Application layer.
- Switches operate on Layer 2 and VLANs are defined on this layer.
- Routers operate on Layer 3 and use ACLs to restrict traffic.

Chapter 3 Practice Questions

1. What protocol does IPv6 use for hardware address resolution?
 - A. ARP
 - B. NDP
 - C. RDP
 - D. SNMP
2. What is the default port for SSH?
 - A. 22
 - B. 23
 - C. 25
 - D. 80
3. You are configuring a host-based firewall so that it will allow SFTP connections. Which of the following is required?
 - A. Allow UDP 21
 - B. Allow TCP 21
 - C. Allow TCP 22
 - D. Allow UDP 22
4. You need to send several large files containing proprietary data to a business partner. Which of the following is the BEST choice for this task?
 - A. FTP
 - B. SNMP
 - C. SFTP
 - D. SSH
5. Your organization is planning to establish a secure link between one of your mail servers and a business partner's mail server. The connection will use the Internet. What protocol is the BEST choice?
 - A. TLS
 - B. SMTP
 - C. HTTP
 - D. SSH
6. You recently learned that a network router has TCP ports 22 and 80 open, but the organization's security policy mandates that these should not be accessible. What should you do?
 - A. Disable the FTP and HTTP services on the router.
 - B. Disable the DNS and HTTPS services on the router.
 - C. Disable the SSH and HTTP services on the router.
 - D. Disable the Telnet and Kerberos services on the router.
7. You need to prevent the use of TFTP through your firewall. Which port would you block?
 - A. TCP 69
 - B. UDP 69
 - C. TCP 21
 - D. UDP 21
8. You need to enable the use of NetBIOS through a firewall. Which ports should you open?
 - A. 137 through 139
 - B. 20 and 21
 - C. 80 and 443
 - D. 22 and 3389

9. Lisa wants to manage and monitor the switches and routers in her network. Which of the following protocols would she use?

- A. Telnet
- B. SSH
- C. SNMP
- D. DNS

10. You need to divide a single Class B IP address range into several ranges. What would you do?

- A. Subnet the Class B IP address range.
- B. Create a virtual LAN.
- C. Create a DMZ.
- D. Implement STP.

11. You need to reboot your DNS server. Of the following choices, which type of server are you MOST likely to reboot?

- A. Unix server
- B. Apache server
- C. BIND server
- D. Web server

12. Your organization is increasing security and wants to prevent attackers from mapping out the IP addresses used on your internal network. Which of the following choices is the BEST option?

- A. Implement subnetting.
- B. Implement secure zone transfers.
- C. Block outgoing traffic on UDP port 53.
- D. Add a WAF.

13. A network technician incorrectly wired switch connections in your organization's network. It effectively disabled the switch as though it was a victim of a denial-of-service attack. What should be done to prevent this in the future?

- A. Install an IDS.
- B. Only use Layer 2 switches.
- C. Install SNMP on the switches.
- D. Implement STP or RSTP.

14. Your organization frequently has guests visiting in various conference rooms throughout the building. These guests need access to the Internet via wall jacks, but should not be able to access internal network resources. Employees need access to both the internal network and the Internet. What would BEST meet this need?

- A. PAT and NAT
- B. DMZ and VPN
- C. VLANs and 802.1x
- D. Routers and Layer 3 switches

15. Your network currently has a dedicated firewall protecting access to a web server. It is currently configured with the following two rules in the ACL along with an implicit allow rule at the end:

PERMIT TCP ANY ANY 443
PERMIT TCP ANY ANY 80

You have detected DNS requests and zone transfer requests coming through the firewall and you need to block them. Which of the following would meet this goal? (Select TWO. Each answer is a full solution.)

- A. Add the following rule to the firewall: DENY TCP ALL ALL 53.
- B. Add the following rule to the firewall: DENY UDP ALL ALL 53.
- C. Add the following rule to the firewall: DENY TCP ALL ALL 25.
- D. Add the following rule to the firewall: DENY IP ALL ALL 53.
- E. Change the implicit allow rule to implicit deny.

16. Your organization wants to prevent users from accessing file sharing web sites. Which of the following choices will meet this need?
- A. Content inspection
 - B. Malware inspection
 - C. URL filter
 - D. Web application firewall
17. Your organization wants to combine some of the security controls used on the network. What could your organization implement to meet this goal?
- A. SSO
 - B. UTM
 - C. VPN
 - D. VLAN
18. Your organization hosts a web server and wants to increase its security. You need to separate all web-facing traffic from internal network traffic. Which of the following provides the BEST solution?
- A. VLAN
 - B. Firewall
 - C. DMZ
 - D. WAF
19. Network administrators connect to a legacy server using Telnet. They want to secure these transmissions using encryption at a lower layer of the OSI model. What could they use?
- A. IPv4
 - B. IPv6
 - C. SSH
 - D. SFTP
20. Which of the following operates on the HIGHEST layer of the OSI model, and is the most effective at blocking application attacks?
- A. IDS
 - B. Router
 - C. WAF
 - D. Stateless firewall

Chapter 3 Practice Question Answers

1. **B.** IPv6 uses the Neighbor Discovery Protocol (NDP) to resolve IPv6 addresses to media access control (MAC) addresses (also called hardware addresses). IPv4 uses the Address Resolution Protocol (ARP) to resolve IPv4 addresses to MAC addresses. Remote Desktop Protocol (RDP) is used to connect to remote systems over port TCP 3389. Administrators use Simple Network Management Protocol (SNMP) to monitor and manage network devices.
2. **A.** Secure Shell (SSH) uses Transmission Control Protocol (TCP) port 22 by default, and it is commonly used with other protocols, such as Secure Copy (SCP) and Secure File Transfer Protocol (SFTP). Telnet uses port 23. SMTP uses port 25. HTTP uses port 80.
3. **C.** You should create a rule to allow traffic using Transmission Control Protocol (TCP) port 22. Secure File Transfer Protocol (SFTP) uses Secure Shell (SSH) on TCP port 22. FTP uses TCP port 21. SSH does not use UDP.
4. **C.** File Transfer Protocol (FTP) is the best choice to send large files, and Secure File Transfer Protocol (SFTP) is the best choice to send large files that need to be protected with encryption. SFTP encrypts data with Secure Shell (SSH) on port 22. FTP data is cleartext and is not suitable for proprietary data. Simple Network Management Protocol (SNMP) is used to manage network devices. Secure Shell (SSH) provides encryption for other protocols, but is not the best choice to send files without combining it with FTP (as SFTP).
5. **A.** Transport Layer Security (TLS) is a good choice to create a secure connection between two systems over the Internet. Although the mails servers will likely exchange mail using Simple Mail Transfer Protocol (SMTP), SMTP by itself will not create a secure link. Similarly, Hypertext Transfer Protocol (HTTP) doesn't create a secure link. Although Secure Shell (SSH) creates a secure connection, it isn't used with SMTP.
6. **C.** You should disable the Secure Shell (SSH) and Hypertext Transfer Protocol (HTTP) services because they use TCP ports 22 and 80 by default. File Transfer Protocol (FTP) uses ports 20 and 21. Domain Name System (DNS) uses port 53. Telnet uses port 23. Kerberos uses port 88.
7. **B.** You should block UDP port 69 to block Trivial File Transfer Protocol (TFTP). TFTP does not use TCP. File Transfer Protocol (FTP) uses TCP port 21.
8. **A.** Network Basic Input/Output System (NetBIOS) uses ports 137 through 139. File Transfer Protocol (FTP) uses ports 20 and 21. Hypertext Transfer Protocol (HTTP) uses port 80 and HTTP Secure (HTTPS) uses port 443. You can connect to remote systems with Secure Shell (SSH) using port 22, and Remote Desktop Protocol (RDP) using port 3389.
9. **C.** Simple Network Management Protocol version 3 (SNMPv3) monitors and manages network devices. She can use Telnet to connect to the devices, but not monitor them. Secure Shell (SSH) is a more secure alternative than Telnet, but it cannot monitor the devices either. Domain Name System (DNS) provides name resolution services.
10. **A.** You can divide any classful IP address range by subnetting it. This breaks up a larger range of IP addresses into smaller network segments or blocks of IP addresses. A virtual local area network (VLAN) divides groups of computers logically, but doesn't use IP ranges. A demilitarized zone (DMZ) is a buffered zone between a protected network and a public network. Spanning Tree Protocol (STP) prevents looping problems caused by incorrect cabling.
11. **C.** Berkeley Internet Name Domain (BIND) is a type of Domain Name System (DNS) software commonly used on the Internet and in some internal networks, so a BIND server is a DNS server. BIND runs on Unix servers, but not all Unix servers are BIND servers. Apache is a type of web server software that runs on Unix and Linux systems.
12. **B.** By implementing secure zone transfers on internal Domain Name System (DNS) servers, it prevents attackers from downloading zone data and mapping out IP addresses and devices. Subnetting divides classful IP address ranges into smaller subnets, but it doesn't prevent attacks. DNS name resolution queries use UDP port 53, so blocking outgoing traffic on UDP port 53 would prevent internal users from using DNS on the Internet. A web application firewall (WAF) protects a web server.
13. **D.** Spanning Tree Protocol (STP) or Rapid STP (RSTP) will prevent switching loop problems. It's rare for a wiring error to take down a switch. However, if two ports on a switch are connected to each other, it creates a switching loop and effectively disables the switch. An intrusion detection system (IDS) will not prevent a switching loop. Layer 2 switches are susceptible to this problem. Administrators use Simple Network Management Protocol (SNMP) to manage and monitor devices, but it doesn't prevent switching loops.
14. **C.** An 802.1x server provides port-based authentication and can authenticate clients. Clients that cannot authenticate (the guests in this scenario) can be redirected to a virtual local area network (VLAN) that grants them Internet access, but not access to the internal network. None of the other solutions provides port security or adequate network separation. Port Address Translation (PAT) and Network Address Translation (NAT) each translate private IP addresses to public IP addresses. A demilitarized zone (DMZ) provides a buffer zone between a public network and a private network for public-facing servers. A virtual private network (VPN) provides access to a private network via a public network. Routers work on Layer 3, and Layer 3 switches mimic some of the functionality of routers.
15. **D, E.** The easiest way is to change the implicit allow rule to implicit deny and that is preferred because it will protect the server from unwanted traffic. You can also deny all IP traffic using port 53 with DENY IP ALL ALL 53. DNS requests use UDP port 53, and zone transfers use TCP port 53 so both UDP 53 and TCP port 53 need to be blocked. You can achieve that goal with DENY IP ALL ALL 53.

16. **C.** A URL filter blocks access to specific web sites based on their URLs. Proxy servers and unified threat management (UTM) devices include URL filters. UTM devices include content inspection to identify and filter out different types of files and traffic, and malware inspection to identify and block malware. A web application firewall (WAF) protects a web server from incoming attacks.
17. **B.** A unified threat management (UTM) device combines multiple security controls into a single device. Single sign-on allows users to sign on once and access multiple resources without signing on again. Users can access a private network over a public network via a virtual private network (VPN). You can configure a virtual local area network (VLAN) on a switch to group computers together logically.
18. **C.** A demilitarized zone (DMZ) is a buffered zone between a private network and the Internet, and it will separate the web server's web-facing traffic from the internal network. You can use a virtual local area network (VLAN) to group computers together based on job function or some other administrative need, but it is created on switches in the internal network. A firewall does provide protection for the web server, but doesn't necessarily separate the web-facing traffic from the internal network. A web application firewall (WAF) protects a web server from incoming attacks, but it does not necessarily separate Internet and internal network traffic.
19. **B.** IPv6 includes the use of Internet Protocol security (IPsec), so it is the best choice and it operates on Layer 3 of the Open Systems Interconnection (OSI) reference model. IPv4 doesn't support IPsec natively. Although you can use Secure Shell (SSH) instead of Telnet, they both operate on Layer 7 of the OSI model. IPv6 operates on Layer 3. Secure File Transfer Protocol (SFTP) is useful for encrypting large files in transit, but it doesn't encrypt Telnet traffic.
20. **C.** A web application firewall (WAF) operates on multiple layers up to Layer 7 of the OSI reference model and blocks attacks against a web server. An intrusion detection system (IDS) also operates on multiple layers up to Layer 7 of the OSI model; however, it is more effective at detecting attacks than blocking them. A router operates on Layer 3 of the OSI model and it can perform packet filtering. A stateless firewall only performs packet filtering and isn't effective against Application layer attacks.

Chapter 4

Securing Your Network

CompTIA Security+ objectives covered in this chapter:

1.1 Implement security configuration parameters on network devices and other technologies.

- VPN concentrators
- NIDS and NIPS (Behavior based, Signature based, Anomaly based, Heuristic)
- Application aware devices (IPS, IDS)

1.2 Given a scenario, use secure network administration principles.

- 802.1x

1.3 Explain network design elements and components.

- Remote Access, Telephony, NAC

1.4 Given a scenario, implement common protocols and services.

- IPSec, TLS, SSL

1.5 Given a scenario, troubleshoot security issues related to wireless networking.

- WPA, WPA2, WEP, EAP, PEAP, LEAP, MAC filter, Disable SSID broadcast, TKIP, CCMP, Antenna Placement, Power level controls, Captive portals, Antenna types, Site surveys, VPN (over open wireless)

2.1 Explain the importance of risk related concepts.

- False positives, False negatives

3.4 Explain types of wireless attacks.

- Rogue access points, Jamming/Interference, Evil twin, War driving, Bluejacking, Bluesnarfing, War chalking, IV attack, Packet sniffing, Near field communication, WEP/WPA attacks, WPS attacks

3.5 Explain types of application attacks.

- Zero-day

3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques.

- Network security (MAC limiting and filtering, 802.1x, Rogue machine detection)
- Security posture (Continuous security monitoring, Remediation)
- Reporting (Alarms, Alerts, Trends)
- Detection controls vs. prevention controls (IDS vs. IPS)

3.7 Given a scenario, use appropriate tools and techniques to discover security threats and vulnerabilities.

- Tools (Honeypots, Honeynets)

4.3 Given a scenario, select the appropriate solution to establish host security.

- Host-based intrusion detection

6.2 Given a scenario, use appropriate cryptographic methods.

- WEP vs. WPA/WPA2 and preshared key
- Use of algorithms/protocols with transport encryption (SSL, TLS, IPSec)

**

In this chapter, you'll learn about some more advanced network security concepts. Topics include intrusion detection systems (IDSs) and intrusion prevention systems (IPSs), methods used to secure wireless networks, and remote access technologies.

Understanding IDSs and IPSs

Intrusion detection systems (IDSs) help detect attacks on systems and networks. Intrusion prevention systems (IPSs) stop attacks in progress by detecting and blocking attacks on systems and networks. Although some active IDSs can also take steps to block attacks, not all IDSs are active. This section presents IDSs first, and then wraps up with some information on IPSs and compares the two.

The primary types of IDSs you'll see are host-based IDSs (HIDSs) and network-based IDSs (NIDSs). Each of these IDSs detects attacks either through predefined attack signatures or by detecting anomalies. Once an attack occurs, an IDS can respond either passively or actively. An IPS responds actively to prevent the attack.

The following items summarize the important concepts related to IDSs and IPSs:

- A HIDS is installed on individual servers and workstations.
- A NIDS is installed on network devices such as routers and firewalls.
- Signature-based (or definition-based) monitoring detects attacks based on known attack patterns.
- Anomaly-based (also called behavior-based or heuristics-based) monitoring detects attacks by first identifying normal operation through a baseline. It then compares current operations against the baseline to detect abnormal behavior.
- A passive IDS logs an alert. It may also inform personnel of the alert.
- An active IDS logs and possibly informs personnel of the alert, and also takes action to change the environment.
- An IPS is similar to an active IDS with one distinctive difference. An IPS is always placed in-line with the traffic so it can prevent the attack from reaching the network.

Chapter 3, “Understanding Basic Network Security,” covers the seven-layer Open Systems Interconnection (OSI) reference model. As a reminder, several security devices such as unified threat management (UTM) devices operate on Layer 7, the Application layer. IDSs and IPSs are also application-aware devices. This allows them to inspect traffic on all layers up to Layer 7.

Packet Sniffing

Chapter 8, “Managing Risk,” discusses protocol analyzers, or sniffers, in more depth, but as an introduction, administrators use them to capture and analyze network traffic sent between hosts. IDSs and IPSs have the same capability. They capture the traffic and analyze it to detect potential attacks or anomalies.

Other tools such as Wireshark can capture transmitted packets over both wired and wireless networks. These tools include additional features that make it easy to analyze the captured packets. Administrators use packet sniffers to troubleshoot issues and attackers use packet sniffers in various attacks.

Sniffing on a wired network requires a physical connection to a network device such as a switch. However, attackers can capture packets sent over a wireless network without a physical connection. For example, someone sitting in a coffee shop that offers free Internet access can easily use a laptop computer to capture all the packets sent over the air by anyone else using the wireless network.

Remember this

Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) include sniffing capabilities.

HIDS

A host-based intrusion detection system (HIDS) is additional software installed on a system such as a workstation or server. It provides protection to the individual host and can detect potential attacks and protect critical operating system files. The primary goal of any IDS is to monitor traffic. For a HIDS, this traffic passes through the network interface card (NIC).

Many host-based IDSs have expanded to monitor application activity on the system. As one example, you can install a HIDS on different Internet-facing servers, such as web servers, mail servers, and database servers. In addition to monitoring the network traffic reaching the servers, the HIDS can also monitor the server applications.

It's worth stressing that a HIDS can help detect malicious software (malware) that traditional antivirus software might miss. Because of this, many organizations install a HIDS on every workstation as an extra layer of protection, in addition to traditional antivirus software. Just as the HIDS on a server is used primarily to monitor network traffic, a workstation HIDS is primarily used to monitor network traffic reaching the workstation. However, a HIDS can also monitor some applications and can protect local resources such as operating system files.

In other organizations, administrators only install a HIDS when there's a perceived need. For example, if an administrator is concerned that a specific server with proprietary data is at increased risk of an attack, the administrator might choose to install a HIDS on this system as an extra layer of protection.

SYN Flood Attack

The SYN flood attack is a common denial-of-service (DoS) attack. Chapter 3 described the three-way handshake to establish a session. As a reminder, one system sends a SYN packet, the second system responds with a SYN/ACK packet, and the first system then completes the handshake with an ACK packet. However, in a SYN flood attack, the attacker sends multiple SYN packets but never completes the third part of the TCP handshake with the last ACK packet.

This is like a friend extending his hand to shake hands with you, you extending your hand in response, and then, at the last instant, the friend pulls his hand away. Although you or I would probably stop extending our hand back to someone doing this, the server doesn't know any better and keeps answering every SYN packet with a SYN/ACK packet.

Each uncompleted session consumes resources on the server, and if the SYN flood attack continues, it can actually crash the server. Some servers reserve a certain number of resources for connections, and once the attack consumes these resources, the system blocks additional connections. Instead of crashing the server, the attack prevents legitimate users from connecting to the server.

IDSs and IPSs can detect a SYN flood attack and respond to block the attack. Additionally, many firewalls include a flood guard that can detect SYN flood attacks and take steps to close the open sessions.

...

NIDS

A network-based intrusion detection system (NIDS) monitors activity on the network. An administrator installs NIDSs sensors on network devices such as routers and firewalls. These sensors gather information and report to a central monitoring server hosting a NIDS console.

A NIDS is not able to detect anomalies on individual systems or workstations unless the anomaly causes a significant difference in network traffic. Additionally, a NIDS is unable to decrypt encrypted traffic. In other words, it can only monitor and assess threats on the network from traffic sent in plaintext or nonencrypted traffic.

Figure 4.1 shows an example of a NIDS configuration. In the figure, sensors are located before the firewall, after the firewall, and on routers. These sensors collect and monitor network traffic on subnets within the network and report to the NIDS console. The NIDS provides overall monitoring and analysis and can detect attacks on the network.




Figure 4.1: NIDS sensors

The decision on where you want to place the sensors depends on what you want to measure. For example, the sensor on the Internet side of the firewall will see all the traffic. However, the sensor on the internal side of the firewall will only see traffic that passes through the firewall. In other words, the firewall will filter some attacks, and the internal sensor won't see them.

If you want to see all attacks on your network, put a sensor on the Internet side. If you only want to see what gets through, put sensors internally only. If you want to see both, put sensors in both places.

Remember this

A HIDS can monitor all traffic on a single host system such as a server or a workstation. In some cases, it can detect malicious activity missed by antivirus software. A network-based IDS (NIDS) is installed on network devices, such as routers or firewalls, to monitor network traffic and detect network-based attacks. A NIDS cannot monitor encrypted traffic and cannot monitor traffic on individual hosts.

Detection Methods

An IDS can only detect an attack. It cannot prevent attacks. In contrast, an IPS prevents attacks by detecting them and stopping them before they reach the target. An attack is any attempt to compromise confidentiality, integrity, or availability.

The two primary methods of detection are signature-based and anomaly-based. Any type of IDS (HIDS or NIDS) can detect attacks based on signatures, anomalies, or both. The HIDS monitors the network traffic reaching its NIC, and the NIDS monitors the traffic on the network.

Signature-Based Detection

Signature-based IDSs (also called definition-based) use a database of known vulnerabilities or known attack patterns. For example, tools are available for an attacker to launch a SYN flood attack on a server by simply entering the IP address of the system to attack. The attack tool then floods the target system with synchronize (SYN) packets, but never completes the three-way Transmission Control Protocol (TCP) handshake with the final acknowledge (ACK) packet. If the attack isn't blocked, it can consume resources on a system and ultimately cause it to crash.

However, this is a known attack with a specific pattern of successive SYN packets from one IP to another IP. The IDS can detect these patterns when the signature database includes the attack definitions. The process is very similar to what antivirus software uses to detect malware. You need to update both IDS signatures and antivirus definitions from the vendor on a regular basis to protect against current threats.

Chapter 8 covers vulnerability scanners in more depth. As an introduction, they can scan systems and networks for vulnerabilities and report issues. They often use vulnerabilities listed in the Common Vulnerabilities and Exposures (CVE) list, which is a dictionary of publicly known security vulnerabilities and exposures. Some IDS systems use the CVE list and its standardized numbering system when identifying and reporting on some issues.

For example, CVE-2014-3522 is a known vulnerability related to Apache servers running on Windows servers. Administrators can look up the number to identify the issue and the resolution. In this case, the resolution is to upgrade Apache to a newer version. You can look up any CVE item at <https://cve.mitre.org/cve/cve.html>.

Anomaly-Based Detection

Anomaly-based (also called heuristic-based or behavior-based) detection first identifies normal operation or normal behavior. It does this by creating a performance baseline under normal operating conditions.

The IDS provides continuous monitoring by constantly comparing current network behavior against the baseline. When the IDS detects abnormal activity (outside normal boundaries as identified in the baseline), it gives an alert indicating a potential attack.

Anomaly-based detection is similar to how heuristic-based antivirus software works. Although the internal methods are different, both examine activity and make decisions that are outside the scope of a signature or definition database.

This can be effective at discovering zero-day exploits. A zero-day vulnerability is usually defined as one that is unknown to the vendor. However, in some usage, administrators define a zero-day exploit as one where the vendor has not released a patch. In other words, the vendor may know about the vulnerability but has not written, tested, and released a patch to close the vulnerability yet.

In both cases, the vulnerability exists and systems are unprotected. If attackers discover the vulnerabilities, they try to exploit them. However, the attack has the potential to create abnormal traffic allowing an anomaly-based system to detect it.

Any time administrators make any significant changes to a system or network that cause the normal behavior to change, they should re-create the baseline. Otherwise, the IDS will constantly alert on what is now normal behavior.

Remember this

Signature-based detection identifies issues based on known attacks or vulnerabilities. Many IDSs use specific numbering systems that they've created or use the CVE list. Signature-based detection systems can detect known anomalies. Anomaly-based IDSs (also called behavior-based) can detect unknown anomalies. They start with a performance baseline of normal behavior and then compare network traffic against this baseline. When traffic differs significantly from the baseline, the IDS sends an alert.

Data Sources and Trends

Any type of IDS will use various raw data sources to collect information on activity. This includes a wide variety of logs, such as firewall logs, system logs, and application logs. These logs can be analyzed to provide insight on trends. These trends can detect a pattern of attacks and provide insight into how to better protect a network.

Many IDSs have the capability to monitor logs in real time. Each time a system records a log entry, the IDS examines the log to determine if it is an item of interest or not. Other IDSs will periodically poll relevant logs and scan new entries looking for items of interest.

Reporting

IDSs report on events of interest based on their settings. All events aren't attacks or actual issues, but instead, they provide a report indicating an event might be an alert or an alarm. Administrators investigate to determine if it is valid. Some systems consider an alarm and an alert as the same thing. Other systems use an alarm for a potentially serious issue, and an alert as a relatively minor issue. The goal in these latter systems is to encourage administrators to give a higher precedence to alarms than alerts.

The actual reporting mechanism varies from system to system and in different organizations. For example, one IDS might write the event into a log as an alarm or alert, and then send an email to an administrator account. In a large network operations center (NOC), the IDS might send an alert to a monitor easily viewable by all personnel in the NOC.

False Positives Versus False Negatives

IDSs are susceptible to both false positives and false negatives. A *false positive* is an alert or alarm on an event that is nonthreatening, benign, or harmless. A *false negative* is when an attacker is actively attacking the network, but the system does not detect it. Neither is desirable, but it's impossible to eliminate both. Most IDSs trigger an alert or alarm when an event exceeds a threshold.

Consider the classic SYN flood attack, where the attacker withholds the third part of the TCP handshake. A host will send a SYN packet and a server will respond with a SYN/ACK packet. However, instead of completing the handshake with an ACK packet, the attacking host never sends the ACK, but continues to send more SYN packets. This leaves the server with open connections that can ultimately disrupt services.

If a system receives one SYN packet without the accompanying ACK packet, is it an attack? Probably not. This can happen during normal operations. If a system receives over 1,000 SYN packets from a single IP address in less than 60 seconds, without the accompanying ACK packet, is it an attack? Absolutely.

With this in mind, administrators set the threshold to a number between 1 and 1,000 to indicate an attack. If administrators set it too low, they will have too many false positives and a high workload as they spend their time chasing ghosts. If they set the threshold too high, actual attacks will get through without administrators knowing about them.

Most administrators want to know if their system is under attack. That's the primary purpose of the IDS. However, an IDS that constantly cries "Wolf!" will be ignored when the real wolf attacks. It's important to set the threshold high enough to reduce the number of false positives, but low enough to alert on any actual attacks.

There is no perfect number for the threshold. Administrators adjust thresholds in different networks based on the network's activity level and their personal preferences.

Remember this

A high incidence of false positives increases the administrator's workload. Administrators often set the IDS threshold high enough that it minimizes false positives but low enough that it does not allow false negatives.

IDS Responses

An IDS will respond after detecting an attack, and the response can be either passive or active. A passive response primarily consists of logging and notifying personnel, whereas an active response also changes the environment to block the attack:

- **Passive IDS.** A passive IDS logs the attack and may also raise an alert to notify someone. Most IDSs are passive by default. The notification can come in many forms, including an email, a text message, a pop-up window, or a notification on a central monitor.
- **Active IDS.** An active IDS logs and notifies personnel just as a passive IDS does, but it can also change the environment to thwart or block the attack. For example, it can modify access control lists (ACLs) on firewalls to block offending traffic, close processes on a system that were caused by the attack, or divert the attack to a safe environment, such as a honeypot or honeynet.

Honeypots

A honeypot is a sweet-looking server—at least it's intended to look sweet to the attacker, similar to how honey looks sweet to a bear. It's actually a server that is left open or appears to have been sloppily locked down, allowing an attacker relatively easy access. The intent is for the server to look like an easy target so that the attacker spends his time in the honeypot instead of in a live network. In short, the honeypot diverts the attacker away from the live network.

As an example, a honeypot could be a web server designed to look like a live web server. It would have bogus data such as files and folders containing fabricated credit card transaction data. If an organization suspects it has a problem with a malicious insider, it can create an internal honeypot with bogus information on proprietary projects.

Honeypots typically have minimal protection that an attacker can easily bypass. If administrators don't use any security, the honeypot may look suspicious to experienced attackers and they may simply avoid it.

Security personnel often use honeypots as a tool to gather intelligence on the attacker. Attackers are constantly modifying their methods to take advantage of different types of attacks. Some sophisticated attackers discover vulnerabilities before a patch is released (also known as a zero-day exploit, or zero-day vulnerability). In some cases, security professionals are able to observe attackers launching zero-day vulnerability attacks against a honeypot.

Honeypots never hold any data that is valuable to the organization. The data may appear to be valuable to an attacker, but its disclosure is harmless. Honeypots have two primary goals:

- **Divert attackers from the live network.** As long as an attacker is spending time in the honeypot, he is not attacking live resources.
- **Allow observation of an attacker.** While an attacker is in the honeypot, security professionals are able to observe the attack and learn from the attacker's methodologies. Honeypots can also help security professionals learn about zero-day exploits, or previously unknown attacks.

Honeynets

A *honeynet* is a group of virtual servers contained within a single physical server, and the servers within this network are honeypots. The honeynet mimics the functionality of a live network.

As an example, you can use a single powerful server with a significant amount of RAM and processing power. This server could host multiple virtual servers, where each virtual server is running an operating system and applications. A physical server hosting six virtual servers will appear as seven systems on a subnet. An attacker looking in will not be able to determine if the servers are physical or virtual.

The purpose of this virtual network is to attract the attention of an attacker, similar to how a single honeypot tries to attract the attention of an attacker. As long as the attacker is in the honeynet, the live network isn't being attacked, and administrators can observe the attacker's actions.

Sun Tzu famously wrote in *The Art of War*, "All warfare is based on deception," and "Know your enemies." Cyberwarfare is occurring daily and security professionals on the front lines of network and system attacks recognize that these attacks mimic warfare in many ways. Honeypots and honeynets provide these professionals with some additional tools to use in this war.

Remember this

Honeypots and honeynets attempt to divert attackers from live networks. They give security personnel an opportunity to observe current methodologies used in attacks, and gather intelligence on these attacks.

Counterattacks

An active response IDS would rarely perform a counterattack against the attacker. Some network security professionals specialize in attacks or counterattacks, but regular administrators should avoid them.

Consider basic human nature. If one person bumps into another in a crowd, the second person could simply ignore it or give a smile and a nod indicating “no problem,” and the event is over. On the other hand, if the response is an aggressive push accompanied by some loud words, the event escalates. It can turn ugly quickly. Now, compare this to some basic facts about attackers today:

- **Attackers are dedicated.** Attackers aren’t just bored teenagers passing their time away like Matthew Broderick in the movie *War Games*. Most attackers today are dedicated criminals working in a semiskilled profession. This is similar to a seasoned car thief with specific skills to break into and steal cars. They are often very good at what they do. Attackers’ skills steadily increase, and their tools are becoming more and more sophisticated.
- **Attackers have unlimited time.** Attackers usually have the luxury of spending 100 percent of their time on attack strategies and methodologies. Compare this with network administrators, who have a host of other duties and rarely can spend 100 percent of their time on security.

Many administrators certainly have the expertise to investigate an attack, trace an IP address back, and launch a counterattack. However, just because you can doesn’t mean you should. The attacker will likely detect the counterattack and escalate the attack. Instead of moving on from your network, the attacker might take your attack personally and consider it a lifelong mission to cripple your network.

It’s also highly likely that the attacking IP address is not the actual attacker. Very often, attackers hijack the machines of unwitting users and launch attacks from their systems. If you counterattack, you could be attacking the wrong computer.

IDS Versus IPS

Intrusion prevention systems (IPSs) are an extension of IDSs. Just as you can have both a HIDS and a NIDS, you can also have a HIPS and a NIPS, but a NIPS is more common. There are two primary distinctions of an IPS when compared with an IDS:

- All IPSs will detect and block attacks, whereas only active IDSs will detect and block attacks. A passive IDS only detects attacks and logs or records them.
- An IPS is in-line with the traffic. In other words, all traffic passes through the IPS and the IPS can block malicious traffic.

As a reminder from the introduction of this section, both IDSs and IPSs have protocol analyzer or sniffer capabilities. This allows them to monitor data streams looking for malicious behavior. An IPS can inspect packets within these data streams and block malicious packets before they enter the network.

In contrast, a NIDS has sensors that monitor and report the traffic. An active NIDS can take steps to block an attack, but only after the attack has started. The in-line configuration of the IPS allows an IPS to prevent attacks from reaching the internal network.

As an example, Figure 4.2 shows the location of two network-based IPSs (NIPS 1 and NIPS 2). All Internet traffic flows through NIPS 1 giving it an opportunity to inspect incoming traffic. NIPS 1 protects the internal network by detecting malicious traffic and preventing attacks from reaching the internal network.




Figure 4.2: NIPS used to detect and prevent attacks

NIPS 2 is protecting an internal private network. As an example, imagine that Homer needs to manage some equipment within a supervisory control and data acquisition (SCADA) network in the nuclear power plant. The SCADA equipment is in the private network. The firewall next to NIPS 2 can have rules that allow traffic from Homer's computer into the network, but block all other traffic. NIPS 2 will then inspect all the incoming traffic and block malicious traffic.

This might seem like overkill, but many advanced persistent threats (APTs) have successfully installed remote access tools (RATs) onto internal systems through phishing or malware attacks. Once the RAT is installed, attackers can now attack from within. If an attacker began launching attacks on the private network from Homer's system, the firewall wouldn't block it. However, the NIPS will prevent this attack from reaching the private network.

Last, notice that each IPS is placed on the edge of the protected network. NIPS 1 is placed on the edge of the network between the Internet and the demilitarized zone (DMZ). NIPS 2 is on the edge of the SCADA network between it and the intranet, or local area network (LAN). This placement ensures that the NIPS can inspect all traffic going into the network.

Remember this

An intrusion prevention system (IPS) is a preventive control. It is similar to an active IDS except that it's placed in-line with traffic. An IPS can actively monitor data streams, detect malicious content, and stop attacks in progress. It can also be used internally to protect private networks.

Securing Wireless Networks

Wireless local area networks (WLANs) have become quite popular in recent years, in both home and business networks. A wireless network is easy to set up and can quickly connect several computers without the need to run cables, which significantly reduces costs.

The significant challenge with wireless networks is security. Wireless security has improved over the years, but wireless networks are still susceptible to vulnerabilities and many users just don't understand how to lock down a wireless network adequately.

A joke by a mainstream comic helps illustrate this. He said, "I think my neighbor must have moved. My Internet connection doesn't work anymore." As with many jokes, it's humorous because it has an element of truth. You might even know some people who have connected to their neighbor's wireless network. Of course, this isn't possible if the owner has secured the wireless network adequately.

Reviewing Wireless Basics

Before digging into wireless security, you need to understand some basic concepts related to wireless devices and networks. If you’ve recently passed the CompTIA Network+ exam, these topics will likely be very familiar to you, but they are still worth looking at to ensure you have an understanding of them from the perspective of the CompTIA Security+ exam.

WAPs and Wireless Routers

A *wireless access point (WAP)* provides wireless clients connectivity to a wired network. Some WAPs connect the wireless clients to the wired network and don’t do anything else. However, many WAPs also have routing capabilities. Vendors market WAPs with routing capabilities as *wireless routers* so that’s how you’ll typically see them advertised. Two distinctions are:

- **All wireless routers are WAPs.** These are WAPs with an extra capability—routing.
- **Not all WAPs are wireless routers.** Many WAPs do not have any additional capabilities. They provide connectivity for wireless clients to a wired network, but do not have routing capabilities.

Figure 4.3 shows a diagram of a wireless router providing connectivity to multiple systems. Notice that the wireless router has both a switch component and a router component, and the drawing at the bottom of Figure 4.3 shows the logical configuration of the network. The devices connect to the switch component and the router component provides connectivity to the Internet through a broadband modem or similar device depending on the Internet Service Provider (ISP) requirements.

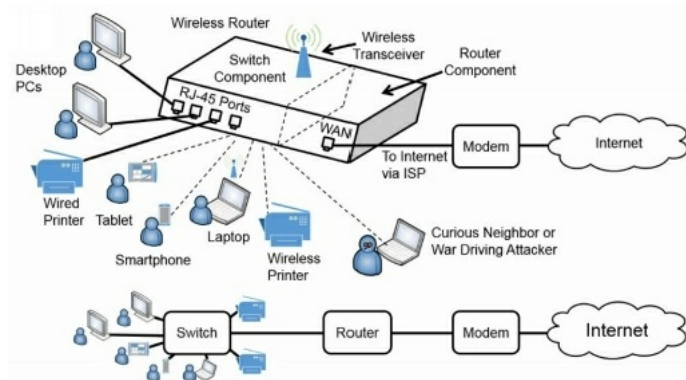


Figure 4.3: Wireless access point with routing capabilities (wireless router)

Most WAPs include physical ports for wired access (labeled as “RJ-45 Ports” in the diagram) and a wireless transceiver for wireless clients. In other words, some users can connect with regular twisted-pair cable, and other users can connect using wireless transmissions. The wired ports and wireless connections all connect through the switch component of the wireless router. Many vendors label the Internet connection WAN for wide area network, but some vendors label this port as “Internet.”

When used as shown in Figure 4.3, the WAP also includes extra services and capabilities such as routing, Network Address Translation (NAT), Dynamic Host Configuration Protocol (DHCP), and more. These extra services reduce the setup time required for the WLAN.

Because wireless networks broadcast on known frequency bands, other wireless users can often see them. This includes both curious neighbors and war driving attackers. War driving is the practice of looking for a wireless network, often by driving around neighborhoods.

802.11

The Institute of Electrical and Electronics Engineers (IEEE) defines many standards, including the IEEE 802.11 group of wireless network protocols. Table 4.1 shows some of the common wireless standards, their maximum throughput, and operating frequencies.

IEEE Standard	Maximum Bandwidth	Operating Frequencies
802.11a	54 Mbit/s	5 GHz
802.11b	11 Mbit/s	2.4 GHz
802.11g	54 Mbit/s	2.4 GHz
802.11n	600 Mbit/s	2.4 GHz and 5 GHz

Table 4.1: Wireless standards

Not every connection to a WAP will achieve the maximum throughput. Instead, the client and WAP negotiate the highest throughput they can achieve without errors. A wireless device in the same room as the WAP will be quicker than a wireless device separated by space, walls, and floors.

Additionally, wireless devices don’t operate on exactly 2.4 GHz or 5 GHz. Instead, wireless transmissions use multiple channels within these bands. For the CompTIA Security+ exam, it’s not critical to know the details of wireless transmission bands, but you should understand that the signals are transmitted over the air, and the frequencies are known. An attacker with a wireless receiver (such as a simple laptop) and a protocol analyzer (a sniffer) can easily detect and capture wireless transmissions.

Antennas

The most commonly used wireless antenna on both WAPs and wireless devices is an omnidirectional (or omni) antenna. *Omnidirectional* antennas transmit and receive signals in all directions at the same time. This allows wireless devices to connect to a WAP from any direction.

Another type of antenna is a directional antenna. A *directional* antenna transmits in a single direction and receives signals back from the same direction. Because the power of the antenna is focused in a single direction, the directional antenna has greater gain than an omni antenna, and it can transmit and receive signals over greater distances.

Although omnidirectional and directional antennas are generic, there are some specific types of antennas used and referenced with wireless devices and networks. They include:

- **Isotropic.** An *isotropic* antenna is a theoretical concept where an antenna has a perfect three-dimensional radiation pattern of 360 degrees vertically and horizontally. In other words, if you measured the power level of the signal 50 feet from any direction, it would have the same power level. Dipole and other omnidirectional antennas attempt to mimic an isotropic antenna.
- **Dipole.** A *dipole* antenna is an actual antenna. Assuming the antenna is standing vertically (such as a pencil standing straight up balanced on the eraser), it has a radiation pattern of 360 degrees horizontally, and about 75 degrees vertically. Some variations of a dipole antenna, such as a folded dipole and half-wave dipole, increase the power gain and overall covered area of the antenna. Most omnidirectional antennas used in wireless networks are a type of dipole antenna.
- **Yagi.** A Yagi (also called Yagi-Uda) antenna is a common type of directional antenna. Yagi antennas typically use a dipole, folded dipole, or half-wave dipole combined with additional elements such as a reflector or director element. These additional elements focus the antenna in a single direction while also increasing the gain and reducing the radiation pattern.

Figure 4.4 provides a generic view of the radiation patterns of different types of antennas. Notice that the omni and dipole antennas don't provide perfect 360-degree coverage. You can also see how the Yagi/directional antenna narrows the radiation pattern while also increasing the gain in a single direction.




Figure 4.4: Directional and omnidirectional antennas

Antenna Power

Wireless access points and antennas include specific measurements to identify their power capabilities. Most devices and antennas include a variation of a decibel (dB) value to express power as a ratio. Three common terms are:

- **dBi.** *Decibels-isotropic* (dBi) identifies the gain of an antenna and is commonly used with omnidirectional antennas. It references an isotropic antenna that can theoretically transmit the signal equally in all directions. Higher numbers indicate the antenna can transmit and receive over greater distances.
- **dBd.** *Decibels-dipole* (dBd) identifies the gain of an antenna compared with a type of dipole antenna. Just as with dBi, higher dBd numbers indicate the antenna can transmit and receive over greater distances.
- **dBm.** *Decibels-milliwatt* (dBm) identifies the power level of the WAP and refers to the power ratio in decibels referenced to one milliwatt. Higher numbers indicate the WAP transmits the signal over a greater distance.

You cannot modify the dBi or dBd gain of an antenna without changing its physical properties. However, many WAPs include power settings allowing you to increase or decrease the power levels. For example, you can decrease the power level to reduce the coverage. Administrators do this occasionally to restrict access to a small area such as a conference room, or to prevent wireless users from connecting from the parking lot or somewhere else outside the building.

Wireless Footprint

The wireless footprint is the area of coverage provided by a WAP or group of WAPs. Figure 4.5 shows a diagram with six normal omnidirectional WAPs (labeled 1 through 6) and two directional WAPs (labeled A and B). The areas outlined with dotted lines indicate the radiation pattern of each of the WAPs, also known as their footprint. Although the omni-WAPs won't have perfect-circle footprints as shown, this does give you an idea of their overall coverage.




Figure 4.5: WAP footprints

As mentioned previously, all WAPs are not wireless routers. In the organization depicted in Figure 4.5, these WAPs would typically not be wireless routers, but instead just plain WAPs. They would provide connectivity for the wireless clients to the wired network.

The two buildings are far enough away from each other and the organization chose to connect the networks using two WAPs with directional Yagi antennas. Because the Yagi antennas provide high gain and a narrow radiation pattern, it reduces the possibility of someone intercepting the signal unless they are directly between the buildings.

Remember this

Most WAPs use an omnidirectional antenna. In some situations, administrators use a high-gain directional Yagi antenna to connect two WAPs together. For example, you can connect two buildings with two WAPs using Yagi antennas.

Notice that the wireless coverage of WAPs 1, 3, 5, and 6 are all uniform. This indicates they have uniform power levels. However, WAP 2 has a smaller footprint, indicating it has a lower power level. In contrast, WAP 4 has a larger footprint, indicating it has a stronger power level. If you want to reduce the footprint of any WAP, you can reduce the power output because the amount of power used by the WAP determines how far it transmits. Use less power and you'll have a weaker signal and a smaller footprint. Of course, the trade-off is reduced performance for authorized users. If the signal is weak, the negotiated speed is slower. Some users farther away from the WAP may not be able to connect at all.

You can also see that there are some dead spots that aren't covered by any WAP. It would be possible to increase the power of all the WAPs to eliminate them. However, this increases the footprint and causes the wireless signal to transmit well beyond the boundaries of the building, which increases the overall risk associated with the wireless network.

Another method of changing the footprint is by modifying the position of the antennas. For example, if you position the antennas vertically (straight up and down), the signals radiate outward, increasing the footprint. However, if you position the antennas horizontally (parallel with the horizon or the floor), the signal radiates up and down more than it radiates outward. This is useful when transmitting a signal between floors of a building, and it also reduces the footprint outside the building.

Administrators have competing goals with the footprint. Users want easy access to the WAP, so users prefer a large footprint with strong signals. However, the stronger the signal is, the easier it is for an attacker to eavesdrop and capture network traffic. From a security perspective, the goal is to limit the footprint to prevent attackers from accessing the wireless network from external locations such as a parking lot, while also ensuring that users have adequate access to the WAP.

Decreasing the footprint isn't always successful at thwarting eavesdroppers. Most common wireless devices use omnidirectional antennas to receive a wireless signal from any direction. However, an attacker can create a directional antenna that can receive wireless traffic from a specific direction. For example, attackers create simple cantennas (antennas using a can) to capture signals from a specific direction. They connect the wireless receiver to one end of an empty can and simply point the can toward a wireless network. By pointing the cantenna in different directions, they can home in on the exact location of a wireless network. Additionally, they can eavesdrop on wireless conversations even though they are well outside the normal footprint.

Remember this

You can limit the range of a WAP to a room or building by reducing the WAP's power level. This prevents people from connecting because they will be out of the WAP's range.

Site Surveys and Antenna Placement

A wireless *site survey* is the process of examining the wireless environment to identify potential issues. Administrators perform a site survey while planning and deploying a WLAN. Administrators or security personnel periodically repeat the site survey to verify the environment hasn't changed and to detect potential security issues.

One method of performing a site survey is to configure a WAP and position the antenna within the organization. Administrators then measure the power levels of the WAP from different areas to determine if it provides the desired coverage. If the WAP doesn't provide adequate coverage, administrators might try to modify the placement of the WAP and/or its antenna, or add additional WAPs.

In addition to measuring power levels within the organization, administrators also measure the power levels outside the organization's perimeter. For example, they would measure the power levels in the parking lot to see if war drivers can access it from there.

Security administrators also perform periodic site surveys as a mitigation and deterrent technique. A site survey can often detect threats such as

rogue access points, evil twins, jamming, and interference. Each of these topics is covered later in this chapter.

Security Protocols

Because wireless networks broadcast over the air, anyone who has a wireless transceiver can intercept the transmissions. You can secure wireless networks with several different steps, but the most important step is to implement a security protocol, such as Wi-Fi Protected Access II (WPA2). The following sections describe the primary security protocols available for wireless networks.

WEP

Wired Equivalent Privacy (WEP) was the original security protocol used to secure wireless networks. As the name implies, the goal was to provide the same level of privacy and security within a wireless network as you'd have in a wired network.

Unfortunately, WEP has significant vulnerabilities, and tools are widely available to break into WEP-protected networks. Due to the widely published vulnerabilities of WEP, the IEEE deprecated the use of WEP in 2004. The IEEE identified WPA as an interim replacement, and WPA2 is a permanent replacement.

WPA

Wi-Fi Protected Access (WPA) was an interim replacement for WEP. It provided an immediate solution to the weaknesses of WEP without requiring users to upgrade their hardware. Even when WPA replaced WEP, its developers recognized that WPA wasn't solid enough to last for an extended period. Instead, WPA improved wireless security by giving users an alternative to WEP with existing hardware while the developers worked on creating the stronger WPA2 protocol.

WPA2

Wi-Fi Protected Access II (WPA2) is the permanent replacement for WEP and WPA. WPA2 (also known as IEEE 802.11i) uses stronger cryptography than both WEP and WPA. The Wi-Fi Alliance requires all devices carrying its WI-FI CERTIFIED logo to meet WPA2 standards, including the use of the Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP).

Some security experts have reported cracking WPA2, though their techniques have not been freely available to the public or reviewed by peers. However, they recommend using a preshared key of at least 20 characters using a complex mix of character types to thwart the vulnerabilities they identified.

Although WPA2 provides significant security improvements over previous wireless encryption techniques, some enterprises need stronger security. Another step you can take is to enable authentication with Enterprise mode, described later in this chapter.

TKIP Versus CCMP

Temporal Key Integrity Protocol (TKIP) is an older encryption protocol used with WPA, and CCMP is a newer encryption protocol used with WPA2. IEEE has deprecated WPA and TKIP due to various security issues, but many wireless networks are still using these older protocols. IEEE recommends using WPA2 with CCMP because it provides significantly more security.

When first released, WPA used Rivest Cipher 4 (RC4) stream encryption with TKIP. WEP also uses RC4. However, TKIP does a better job of managing the encryption keys than WEP does, making it more secure. Additionally, TKIP encrypts each packet with a new key. Even though TKIP corrects several of WEP's flaws, it was ultimately cracked.

A benefit of TKIP is that it didn't require new hardware. WEP users could upgrade software and/or firmware and implement WPA with TKIP without the need to replace the hardware. Newer hardware supports WPA2, so the usage of WPA and TKIP is waning. However, you may still see some legacy hardware using WEP, WPA, and TKIP.

Later implementations of WPA support Advanced Encryption Standard (AES) instead of TKIP. Chapter 10, "Understanding Cryptography," presents AES in more depth, but in short, it is a very strong and efficient encryption algorithm. Many applications beyond WPA/WPA2 use AES to provide secure encryption and ensure confidentiality. Several people have been successful at cracking WPA with TKIP, so whenever possible, it's best to upgrade WPA to WPA2, or at least upgrade TKIP to use AES.

WPA2 supports CCMP, which is based on AES and is much stronger than WPA using TKIP. WPA2 also employs much more secure methods of managing the encryption keys than either WEP or WPA.

Remember this

WEP has several weaknesses and should not be used. WPA provided an immediate replacement for WEP and originally used TKIP with RC4, which was compatible with older hardware. Later implementations support the stronger AES encryption algorithm. WPA2 is the permanent replacement for WEP and WPA. WPA2 supports CCMP (based on AES), which is much stronger than the older TKIP protocol.

IEEE 802.1x

Chapter 3 introduced the 802.1x protocol in the context of port security. As a reminder, an 802.1x server is integrated with a database of accounts and it provides port-based authentication by requiring users and devices to authenticate before granting them access to a network. When systems connect, the 802.1x server challenges them to authenticate and prevents full network access until it receives valid credentials. This prevents rogue devices from being able to access network resources.

You can implement IEEE 802.1x as a Remote Authentication Dial-In User Service (RADIUS) server. RADIUS (explained in Chapter 1, “Mastering Security Basics”) provides centralized authentication. When implemented with WPA or WPA2, 802.1x provides an added layer of protection by ensuring users can authenticate before granting them access to the wireless network.

At some point, people started saying that 802.1x is shorthand for multiple wireless protocols such as 802.11a, 802.11b, and so on. It is not. You *can* implement 802.1x with WPA and WPA2 using Enterprise mode, described in the next section.

Personal Versus Enterprise Modes

Both WPA and WPA2 can operate in either Personal or Enterprise modes. When using Personal mode, users access the wireless network anonymously with a preshared key (PSK) or passphrase. This doesn’t provide authentication. As a reminder, authentication (presented in Chapter 1) proves a user’s identity with the use of credentials such as a username and password. Users claim an identity with a username and prove their identity with a password.

In contrast, WPA or WPA2 Enterprise mode forces users to authenticate with unique credentials before granting them access to the wireless network. Enterprise mode uses an 802.1x server, often implemented as a RADIUS server, which accesses a database of accounts. If users don’t have the proper credentials, Enterprise mode (using an 802.1x server) blocks their access. Also, an 802.1x server has a certificate on it to secure the authentication process.

Figure 4.6 shows two screenshots of a Cisco wireless router with the Wireless Security section selected. By clicking in the box next to Security Mode, you can select a variety of different security modes such as WEP, WPA Personal, WPA2 Personal, WPA Enterprise, or WPA2 Enterprise. When you select one of the Personal settings such as WPA2 Personal in the top portion of the figure, it shows a passphrase. It can be as many as 63 characters long and the passphrase you enter here is the same passphrase you would enter on all the wireless devices. Many security experts recommend using a passphrase at least 20 characters long, with a mix of uppercase, lowercase, numbers, and special characters.

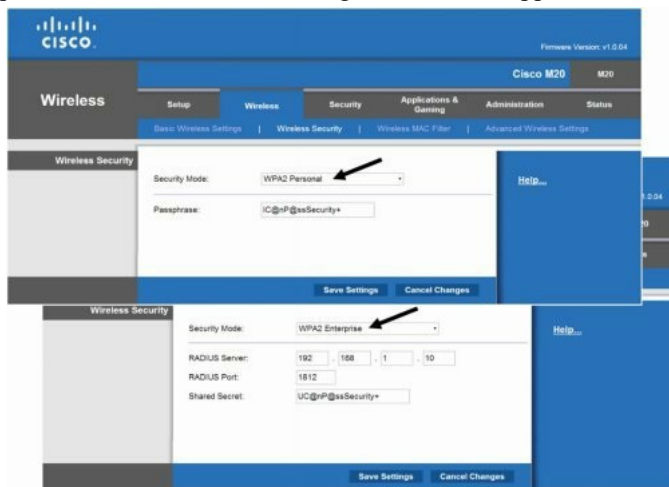


Figure 4.6: Configuring wireless security

If you select WPA2 Enterprise, as shown in the bottom portion of the figure, it displays different information. You would need to put in the IP address of the RADIUS (or 802.1x) server, the port it is using, and a shared secret. The official default port for RADIUS is 1812. However, some vendors have used other ports such as 1645. The key is that you must enter the same port here that the server is using. The shared secret is similar to a password and you must enter it here exactly as it is entered on the server.

After configuring WPA2 Enterprise on a WAP, it redirects all attempts to connect to the RADIUS server to authenticate. After users authenticate, the RADIUS server tells the WAP to grant them access.

Wireless authentication systems using an 802.1x server are more advanced than most home networks need, but many larger organizations use them. In other words, most home networks use Personal mode, whereas many organizations use Enterprise mode to increase security. A combination of both a security protocol such as WPA2 and an 802.1x authentication server significantly reduces the chance of a successful access attack against a wireless system.

Remember this

Personal mode (or WPA-PSK and WPA2-PSK) uses a preshared key and does not provide individual authentication. WPA/WPA2 Enterprise mode is more secure than Personal mode, and it provides strong

authentication. Enterprise mode uses an 802.1x server (implemented as a RADIUS server) to add authentication.

EAP, PEAP, and LEAP

The Extensible Authentication Protocol (EAP) is an authentication framework that provides general guidance for authentication methods. 802.1x servers typically use one of these methods to increase the level of security during the authentication process. Some methods are:

- **EAP.** EAP provides a method for two systems to create a secure encryption key, also known as a Pairwise Master Key (PMK). Systems then use this key to encrypt all data transmitted between the devices. Both TKIP or AES-based CCMP use this key, though CCMP is much more secure.
- **Protected EAP (PEAP).** PEAP provides an extra layer of protection for EAP. The EAP designers assumed that EAP would be used with adequate physical security to ensure the communication channel was secure. In practice, that wasn't always the case, but PEAP protects the channel. PEAP encapsulates and encrypts the EAP conversation in a Transport Layer Security (TLS) tunnel. PEAP requires a certificate on the server, but not the clients. A common implementation is with Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2).
- **EAP-Tunneled TLS (EAP-TTLS).** This is an extension of PEAP, allowing systems to use some older authentication methods such as Password Authentication Protocol (PAP) within a TLS tunnel. It requires a certificate on the 802.1x server but not the clients.
- **EAP-TLS.** This is one of the most secure EAP standards and is widely implemented. The primary difference between PEAP and EAP-TLS is that it requires certificates on the 802.1x server and on each of the wireless clients.
- **Lightweight EAP (LEAP).** Cisco created LEAP using a modified version of the Challenge Handshake Authentication Protocol (CHAP) presented in Chapter 1. LEAP does not require a digital certificate. Most wireless devices support LEAP, but it is susceptible to a known attack. Cisco recommends using stronger protocols, instead of LEAP.

Note that PEAP, EAP-TTLS, and EAP-TLS all use digital certificates. Chapter 10 digs into certificates much deeper, but as an introduction, certificates help provide strong authentication and encryption services. However, a certificate authority (CA) must issue certificates, so an organization must either purchase certificates from a public CA, or implement a private CA within the network.

Remember this

Enterprise mode requires an 802.1x server. PEAP and EAP-TTLS require a certificate on the 802.1x server. EAP-TLS also uses TLS, but it requires certificates on both the 802.1x server and each of the clients. LEAP is proprietary to Cisco.

WTLS and ECC

Two other security protocols you may run across are Wireless Transport Layer Security (WTLS) and elliptic curve cryptography (ECC). Many smaller wireless devices use WTLS or ECC.

Smaller wireless devices such as PDAs and cell phones don't have the same processing power as servers and desktop computers and can't easily handle the processing requirements of advanced security protocols such as WPA2.

However, WTLS and ECC provide protection for the smaller devices without requiring a significant amount of processing power. WTLS is a wireless implementation of TLS. ECC elegantly exploits a field of mathematics that can use a formula to create a curve and another formula to identify one or more points on the curve. ECC uses the points on the curve to create the encryption key.

What Wireless Security Are You Using?

What wireless security are you using? WEP, WPA, WPA2, or nothing at all?

I frequently ask this question when teaching CompTIA Security+ classes. Many students simply don't know, and this question often starts a lot of whispering among students, as some realize they may not even be using WEP.

The next day, many of the students report their findings, and it's common to hear that they've upgraded security on their wireless networks. You can upgrade many WAPs and wireless NICs to support WPA and WPA2 if they don't already support it. However, hardware on some older wireless NICs will only support WPA.

How about you? What wireless security are you using? If you don't know, check. The single most important step you can take to secure your wireless network is to upgrade to WPA2. Then ensure the PSK is at least 20 characters long with a complex mix of uppercase, lowercase, numbers, and characters. Go ahead. Dig out the manual for your wireless access point or wireless router and check it out now.

. . .

Captive Portals

A *captive portal* is a technical solution that forces clients using web browsers to complete a specific process before it allows them access to the network. Organizations commonly use it as a hot spot that requires users to log on or agree to specific terms before they can access the Internet. Here are three common examples:

- **Free Internet access.** Many hospitals and other medical facilities provide free Internet access to patients and visitors. The captive portal requires users to acknowledge and agree to abide by an acceptable use policy. Free captive portals rarely require users to log on, but instead just require them to check a box indicating they agree, and then click a button to continue.
- **Paid Internet access.** Many hotels, resorts, cruise ships, and airlines provide Internet access to customers, but on a pay-as-you-go basis. When users attempt to access the Internet, they are redirected to the captive portal and must successfully log on with a precreated account or enter credit card information to pay for access.
- **Alternative to IEEE 802.1x.** Adding an 802.1x server can be expensive and is sometimes not a feasible option. Organizations can use captive portals as an alternative. It requires users to authenticate before granting them access.

Hot Spots and Isolation Mode

Many small business owners create wireless hot spots for their customers and you might be asked to help them configure one. If you want to prevent wireless clients from communicating with each other, you can enable Isolation mode on the WAP. Clients are able to connect to the WAP, but Isolation mode segments or isolates each wireless user. This provides a level of security for the customers, but does not prevent someone from hosting an evil twin.

Another consideration is the type of security. Many hot spots use WEP with Open System Authentication (OSA), which doesn't use a preshared key, or they just disable the security. When configured this way, hot spot administrators often use a captive portal to provide a warning to the users that their communications are not secure.

Remember this

Isolation mode is used in an access point (AP) to prevent clients from connecting to each other. Public networks sometimes use this to protect wireless clients. You can configure hot spots with WEP and Open System Authentication or security disabled so that users do not need a preshared key.

Other Security Concerns

The use of WPA2, and especially WPA2 Enterprise, clearly provides the highest level of security for wireless networks. However, you can take some additional steps to secure them. The settings described in this section are normally accessible via a group of web pages hosted on your wireless router. You can often access these web pages with your web browser by entering either *http://192.168.0.1* or *http://192.168.1.1* to access the home page.

The Configuring a Wireless Router Lab shows how to configure several security settings on a wireless router. Although your wireless router might be a little different, you'll still be able to see many of the typical configuration settings. You can access this lab and other online exercises for this book at *http://gcgapremium.com/labs/*.

Change Default Administrator Password

Many WAPs come with a default Administrator account of “admin,” and default passwords of “admin.” Some even ship with blank passwords. The WAP's technical manual identifies the default account names and passwords and most manuals stress changing the password. However, many home users do not change the default password.

If the default password isn't changed, anyone who can access your WAP can log on and modify the configuration. Additionally, anyone with access to the Internet can easily download instruction manuals for the popular WAPs to identify the default administrator names and passwords. As an example, *http://portforward.com/* has lists of usernames and passwords for a wide assortment of routers.

An attacker can easily bypass an otherwise secure wireless network if the administrator password is not changed. The attacker can log on and simply turn off security. Unless you go back into the WAP configuration, you may never know that security is disabled.

Enable MAC Filtering

An additional step you can take to provide a small measure of security to a wireless network is to enable media access control (MAC) filtering. As a reminder, the MAC address (also called a physical address or hardware address) is a 48-bit address used to identify network interface cards (NICs). You will usually see the MAC address displayed as six pairs of hexadecimal characters such as 00-16-EA-DD-A6-60. Every network interface card (NIC) including wireless adapters has a MAC address.

MAC filtering is a form of network access control. It's used with port security on switches (covered in Chapter 3), and you can use it to restrict access to wireless networks.

For example, Figure 4.7 shows the MAC filter on a Cisco WAP. In the figure, you can see that the system is set to Permit PCs Listed Below to Access the Wireless Network. The MAC 01 through MAC 02 text boxes include MAC addresses of two devices.

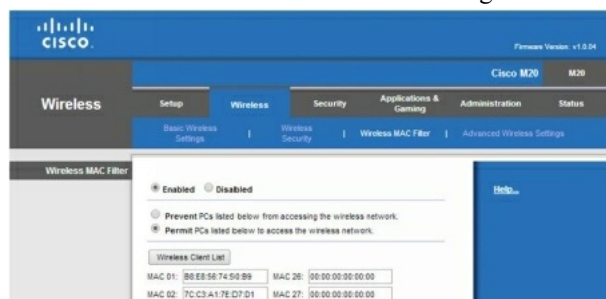


Figure 4.7: MAC filter on a WAP

Theoretically, MAC addresses are unique. With this in mind, the MAC filter in Figure 4.7 limits access to only the two devices with these MAC addresses. This may sound secure, but an attacker with a wireless sniffer can easily identify the MAC addresses allowed in a wireless network. Additionally, it's very easy to change a MAC address. An attacker can launch a spoofing attack by changing the MAC address on his laptop to impersonate one of the allowed MAC addresses.

Many operating systems include built-in functionality to change a NIC's MAC address. For example, in Windows 7 you can access the NIC's properties from Device Manager, click the Advanced tab, and configure the Network Address setting with a new MAC.

Remember this

MAC filtering can restrict access to a wireless network to specific clients. However, an attacker can use a sniffer to discover allowed MAC addresses and circumvent this form of network access control. It's relatively simple for an attacker to spoof a MAC address.

War Driving

War driving is the practice of looking for a wireless network. Although war driving is more common in cars, you can just as easily do it by walking around in a large city. Attackers use war driving to discover wireless networks that they can exploit and often use directional antennas (cantennas) to detect wireless networks with weak signals.

Administrators sometimes use war driving as part of a wireless audit. A wireless audit is a detective control and examines the signal footprint, antenna placement, and encryption of wireless traffic. These audits are useful at detecting weaknesses in wireless networks. For example, administrators can sometimes detect the existence of rogue access points and evil twins by war driving, and determine when their WAP's footprint extends too far.

Many current operating systems include software to identify wireless networks. For example, Microsoft Windows 7 includes tools that allow you to view details about wireless networks in range of the system. The software shows the name of the wireless network, its signal strength, and the security protocol used (such as WEP, WPA, or WPA2).

Remember this

Administrators use war driving techniques as part of a wireless audit. A wireless audit checks a wireless signal footprint, power levels, antenna placement, and encryption of wireless traffic. Wireless audits using war driving can detect rogue access points and identify unauthorized users.

War Biking

In April 2014, James Lyne, global head of security research at Sophos, went “war biking” in London. He attached a 4G modem onto his bike and configured it with a Wi-Fi hot spot, complete with portal pages offering free Internet access. He then rode around the city.

In just two days, over 2,900 users connected to his hot spot. Many of them accessed web sites requiring usernames and passwords, such as online banking sites. Although he mentioned it would have been trivially easy to collect user data, including usernames and passwords, he didn't do so. He routinely destroys all personal user logon data he collects, but a criminal using the same methodologies won't destroy the data. In a similar experiment in San Francisco, over 1,500 users connected to his war bike, and in Las Vegas, about 4,000 users signed on.

During the tests, Lyne also mimicked basic war driving techniques looking for wireless networks. He found close to 30 percent of the networks using either no encryption or WEP, and about another 30 percent using WPS. Less than 20 percent of the networks used the secure WPA2 protocol.

Although he could have configured his wireless network as an evil twin, using the same SSID as legitimate networks, he didn't do so. He found that the majority of users connect first and “ask questions later.” He compared their behavior as being similar to “shouting out personal bank account numbers and passwords at a café and being shocked that other people steal them.” You can learn more about Lyne's war biking tours here:

<http://www.sophos.com/warbiking>.

War Chalking

Occasionally, people use war chalking to publicly mark wireless networks. These are simple marks written in chalk, or painted on a wall as graffiti. For example, two parentheses marks placed back to back as)(indicate an open Wi-Fi network, and an open circle with a W in the middle indicates a Wi-Fi network using WEP. This isn't very common anymore, primarily because businesses often have signs indicating they have wireless capabilities for their customers.

Change Default SSID

Wireless networks are identified by a service set identifier (SSID), which is simply the name of the wireless network. Many WAPs come with default SSIDs. For example, the default SSID of some older Linksys WAPs is “Linksys” or “linksys-g.” Some newer WAPs force you to enter the name of the SSID when you first install it and do not include a default. From a defense-in-depth perspective, it's a good idea to change the name of the SSID if a default is used. It simply gives attackers less information.

For example, if a war driver sees a wireless network with a SSID of Linksys, the attacker has a good idea that the network is using a Linksys WAP. If the attacker knows about specific weaknesses with this WAP, he can start exploiting these weaknesses. On the other hand, a WAP with a SSID of “Success” doesn't give the attacker any clues about the WAP.

Disable SSID Broadcasting or Not

One of the goals of 802.11 wireless networks is ease of use. The designers wanted wireless computers to be able to easily find each other and work together. They were successful with this goal. Unfortunately, attackers can also easily find your networks. By default, WAPs broadcast the SSID in cleartext, making it easy to locate wireless networks.

At some point years ago, someone stated that the SSID was a password (not true!), and many information technology (IT) professionals latched onto

the idea that you can increase security by disabling the SSID broadcast. Others say that the SSID has nothing to do with security and disabling the broadcast reduces usability but does not increase security.

As background, WAPs must regularly send out a beacon frame to ensure interoperability with other devices in the wireless network. This beacon frame includes the SSID, and if the SSID broadcast is disabled, the SSID entry is blank. However, even if the SSID broadcast is disabled, the WAP includes the SSID in Probe responses sent in response to Probe requests from authorized wireless clients. Because of this, it's easy for an attacker with a wireless sniffer to listen for the Probe responses and detect the SSID.

In other words, disabling the SSID makes it a little more difficult for attackers to find your network, but not much. It's almost like locking the front door of your house, but leaving the key in the lock.

Steve Riley wrote in a security blog titled "Myth vs. reality: Wireless SSIDs" that disabling the SSID for security "is a myth that needs to be forcibly dragged out behind the woodshed, strangled until it wheezes its last labored breath, then shot several times for good measure." In case it isn't clear, Mr. Riley is in the camp that says you should not disable the SSID. For the record, I agree with him.

For the CompTIA Security+ exam, you should know that it is possible to disable the SSID broadcast and hide the network from casual users. However, an attacker with a wireless sniffer can easily discover the SSID even if SSID broadcast is disabled.

Remember this

The service set identifier (SSID) identifies the name of the wireless network. You should change the SSID from the default name. Disabling SSID broadcast can hide the network from casual users, but an attacker can easily discover it with a wireless sniffer.

WEP/WPA Attacks

Chapter 10 covers encryption in greater depth, but an important concept is that most encryption methods use both an encryption algorithm and a key. Encryption algorithms are public and stay constant. However, the keys are secret and change regularly with symmetric encryption algorithms such as RC4.

As an example, imagine a key of A1B2C3 hexadecimal. The sending system uses RC4 with this key to encrypt data before sending it. The receiving system then uses this key to decrypt the data. When sending other data, the systems would use a different key (such as C4B3A2 hexadecimal) to encrypt and decrypt the data.

WEP uses the RC4 stream cipher for encryption of transmitted data. An important implementation rule with any stream cipher is "do not reuse encryption keys." WEP broke this implementation rule making it easy for attackers to discover encryption keys. Once they discover an encryption key, they are also able to discover the WEP key or passphrase, and decrypt all data.

IV Attacks

The encryption keys used in WEP are derived from the WEP key entered into the WAP and the wireless devices. WEP then uses an initialization vector (IV) to create the keys used for RC4. The IV used by WEP is a relatively small 24-bit number.

As an example, imagine the passphrase is Ip@\$\$ed. WEP combines Ip@\$\$ed with an IV to create a key, and then encrypts the packet with this key. Of course, the receiving system needs to be able to decrypt the packet. WEP includes the IV in plaintext. Both entities already know the passphrase, so it isn't included. When the other system receives the packet, it combines the IV included in the packet with the passphrase to decrypt the packet.

Although RC4 is very strong and used successfully elsewhere, the small IV used in WEP resulted in systems reusing keys, and made it quite easy for attackers to crack it. Attackers have off-the-shelf software they can use to capture, analyze, and decrypt WEP traffic.

In many IV attacks, the attacker uses packet injection techniques to add additional packets into the data stream. The WAP responds with more packets, increasing the probability that it will reuse a key. An IV attack using packet injection decreases the time it takes to crack a WEP key to a very short time, sometimes less than a minute.

Remember this

WEP is weak and should not be used. It has several security problems, including the use of weak IVs used to create encryption keys for the otherwise secure RC4 symmetric encryption protocol. In an IV attack, the attacker uses packet injection, increasing the number of packets to analyze, and discovers the encryption key.

WPA Cracking Attacks

Unfortunately, WPA is susceptible to password-cracking attacks, especially when the network is using a weak PSK or passphrase. A WPA cracking attack involves three steps:

1. **Use a wireless sniffer or protocol analyzer to capture wireless packets.** Sniffing this way is similar to how attackers sniff wired networks to eavesdrop and capture information sent across a network.
2. **Wait for a wireless client to authenticate.** WPA wireless clients authenticate with WAPs using a four-way handshake where they

exchange information. Essentially, the client needs to prove to the WAP that it knows the passphrase. However, the client doesn't send the passphrase in cleartext. Instead, the four-way handshake allows the client to encrypt the passphrase in such a way that the WAP can decrypt it and verify that the client has the correct passphrase.

3. **Use a brute force attack.** Once attackers have the encrypted passphrase from the captured four-way handshake, they can launch an offline brute force attack. Automated tools such as Aircrack-ng compare the encrypted password in the capture against passwords in one or more password files. When successful, it gives the attacker the actual passphrase used by the WLAN.

A key here is that the attacker must capture the four-way handshake. This only occurs when a user tries to authenticate. If no one is active on the network, this attack won't work. However, if a user is active, the attacker can use a different attack, such as a jamming attack, to disconnect the user and force the user's system to authenticate again.

Remember this

In WPA cracking attacks, attackers capture traffic with a wireless sniffer waiting for an authorized client to connect, so that they can capture the four-way authentication handshake information. They then use a brute force attack to discover the passphrase.

WPS Attacks

Wi-Fi Protected Setup (WPS) allows users to configure a wireless network without typing in the passphrase. Instead, users can configure devices by pressing buttons or by entering a short personal identification number (PIN).

For example, a user can configure a new wireless device by pressing a button on the WAP and on the wireless device. It will automatically configure the device within about 30 seconds with no other actions needed. These buttons can be physical buttons on the devices, or virtual buttons that the user clicks via an application or web page. When using the PIN method, users first identify the eight-digit PIN on the WAP and then enter the PIN on the new wireless device.

Unfortunately, WPS is susceptible to brute force attacks. A WPS brute force attack keeps trying different PINs until it succeeds. Reaver is an open source tool freely available that allows attackers to discover the PIN within 10 hours, and often much quicker. Once it discovers the PIN, it can then discover the passphrase in both WPA and WPA2 wireless networks.

Security experts recommend disabling WPS on all devices. However, not all devices include the capability to turn off WPS. Worse, many WAP interfaces include configuration settings that appear to turn off WPS—making users think it's disabled when it's still operational and vulnerable to attacks. Several testers reported that they were unable to disable WPS on each Linksys and Cisco Valet WAPs they tested. Some vendors have released firmware updates to address this, but updates are not available for all devices.

Rogue Access Points

Generically, you can think of a rogue as a scoundrel, a crook, or a villain. A rogue access point is a WAP placed within a network by someone with some type of attack in mind. Clearly, if a rogue is a crook or villain, then rogue access points are not an administrator's friend. You may also see them called counterfeit access points, which is also a clear indication they aren't legitimate.

Attackers may connect a rogue access point to network devices in wireless closets that lack adequate physical security. This access point acts as a sniffer to capture traffic passing through the wired network device, and then broadcasts the traffic using the wireless capability of the WAP. The attacker can then capture the exfiltrated data files while sitting in the parking lot. Data exfiltration is the unauthorized transfer of data from an organization to a location controlled by an attacker.

Additionally, attackers may be able to use the rogue access point to connect into the wired network. This works the same way that regular users can connect to a wired network via a wireless network. The difference is that the attacker configures all the security for the counterfeit access point and can use it for malicious purposes.

If you discover an unauthorized WAP, you should disconnect it as quickly as possible. A basic first step to take when you discover any attack is to contain or isolate the threat. By simply unplugging an Ethernet cable, you can stop any attacks from an unauthorized WAP.

Often, administrators will use war driving tools to scan their networks for rogue access points. This can help identify the physical location of access points because the signal will get stronger as the administrator gets closer. Some sophisticated war driving tools include directional antennas (such as cantennas) that an administrator (or an attacker) can use to locate a WAP.

Evil Twins

An evil twin is a rogue access point with the same SSID as a legitimate access point. For example, many public places such as coffee shops, hotels, and airports include free Wi-Fi as a service. An attacker can set up a WAP using the same SSID as the public Wi-Fi network, and many unsuspecting users will connect to this evil twin.

Once a user connects to an evil twin, wireless traffic goes through the evil twin instead of the legitimate WAP. Often, the attacker presents bogus logon pages to users in an attempt to capture usernames and passwords. Other times, they simply capture traffic from the connection, such as email or text entered into a web page, and analyze it to detect sensitive information they can exploit.

Although it might sound complex to set up an evil twin, it's actually rather easy. Attackers can configure a laptop that has a wireless access card as a WAP. With it running, the attackers look just like any other user in a coffee shop or airport waiting area. They'll have their laptop open and appear to be working (just like you perhaps), and you'll have no idea they are trying to steal your credentials or other personal data that you send over the Internet via the evil twin. Similarly, attackers can set one up in a parking lot or another location close to an organization and try to trick employees or visitors.

Periodic site surveys or audits can often detect the presence of both rogue access points and evil twins. As an example, consider Table 4.2, which shows the result of a wireless site survey. It indicates measurements for access points from different locations in the organization.

MAC	SSID	Encryption	Power
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	47
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	62
12:AB:56:EF:34:CD	GetCertifiedGetAhead	WPA2	20
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	57
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	49

Table 4.2: Wireless site survey results

Notice that the SSID is the same for all five measurements and they are all using the same encryption. However, the MAC address of the third entry is different from the others. This indicates it is a different AP. Also, the power level is significantly lower than the others, so it is probably not operating from within the organization, but might be operating from a car in the parking lot.

Remember this

Rogue access points are malicious and often used to capture and exfiltrate data. An evil twin is a rogue access point using the same SSID as a legitimate access point. Whereas a secure WAP blocks unauthorized users, a rogue access point provides access to unauthorized users.

Jamming and Interference

Attackers can transmit noise or another radio signal on the same frequency used by a wireless network. This interferes with the wireless transmissions and can seriously degrade performance. This type of denial-of-service attack is commonly called jamming and it usually prevents all users from connecting to a wireless network. In some cases, users have intermittent connectivity because the interference causes them to lose their association with the WAP and forces them to try to reconnect.

In some cases, you can increase the power levels of the WAP to overcome the attack. However, it's worth remembering that as you increase the power level, you increase the wireless footprint and become more susceptible to war driving attacks.

Another method of overcoming the attack is to use different wireless channels. Each wireless standard has several channels you can use, and if one channel is too noisy, you can use another one. Although this is useful to overcome interference in home networks, it won't be as effective to combat an interference attack. If you switch channels, the attacker can also switch channels.

Near Field Communication Attacks

Near field communication (NFC) is a group of standards used on mobile devices that allow them to communicate with other mobile devices when they are close to them. For example, in some cases, you can share information with a friend just by touching her smartphone with yours, or by placing your smartphone in close proximity to hers.

Many credit card readers support payments using NFC technologies. For example, you can make a purchase in some fast-food restaurants just by placing your phone close to the credit reader. These NFC technologies use similar technologies as proximity card readers discussed in Chapter 2, “Exploring Control Types and Methods.”

Bluetooth Wireless

Bluetooth is a short-range wireless system used in personal area networks (PANs) or a network of devices close to a single person. Bluetooth devices include smartphones, personal digital assistants (PDAs), and computer devices.

The range of Bluetooth was originally designed for about three meters (about 10 feet), but the range is often farther, and ultimately extends beyond a person's personal space. Attackers have found that attacks on these networks are possible. Two common attacks are bluesnarfing and bluejacking. Both attacks are much easier when Bluetooth devices remain in Discovery mode.

Discovery Mode

When Bluetooth devices are first configured, they are configured in Discovery mode. Bluetooth devices use MAC addresses, and in Discovery mode the Bluetooth device broadcasts its MAC address, allowing other devices to see it and connect to it. This is required when pairing Bluetooth devices.

For example, if you had a cell phone and an earpiece that both supported Bluetooth, you would use Discovery mode to pair the two devices. You could then keep the cell phone in your pocket or purse, but still carry on a conversation through the earpiece.

One of the risks with Bluetooth occurs when a Bluetooth device is left in Discovery mode. Just as you can pair an earpiece with your cell phone, an attacker can pair a Bluetooth-enabled laptop with your cell phone if it's left in Discovery mode. With the right software, the attacker can then launch bluesnarfing and bluejacking attacks (described in the next section).

Although there are improvements in Bluetooth devices, such as using PINs or passwords, it's still important to ensure that Discovery mode is disabled after pairing them. With Discovery mode disabled, the device doesn't broadcast information about itself. Additionally, many devices add encryption to the communication process when Discovery mode is disabled.

Bluetooth Attacks

Several attack methods target Bluetooth devices specifically. These include:

- **Bluejacking.** This is the practice of sending unsolicited messages to nearby Bluetooth devices. Bluejacking messages are typically text, but can also be images or sounds. Bluejacking is relatively harmless, but does cause some confusion when users start receiving messages.
- **Bluesnarfing.** Any unauthorized access to or theft of information from a Bluetooth connection is bluesnarfing. A bluesnarfing attack can access information, such as email, contact lists, calendars, and text messages. Attackers use tools such as hcitool and obexftp.
- **Bluebugging.** Bluebugging attacks allow an attacker to take over a mobile phone. Attackers can listen in on phone conversations, enable call forwarding, send messages, and more.

The single best protection against all these attacks is to ensure that Bluetooth devices are not left in Discovery mode.

Remember this

Bluesnarfing is the unauthorized access to or theft of information from a Bluetooth device. *Bluejacking* is the unauthorized sending of text messages to a nearby Bluetooth device. Ensuring devices are not left in Discovery mode is a primary protection for Bluetooth devices.

Exploring Remote Access

Remote access is the group of technologies that allow users to access an internal network from remote locations. Remote Access Service (RAS) provides access through dial-up or virtual private networks (VPNs). Several components come together to form a successful RAS solution. They include access methods such as through dial-up or a VPN, authentication methods, and network access control (NAC) technologies to inspect clients before granting them access to the network.

Chapter 1 covers authentication topics in depth and includes information on authentication services. This section covers access methods and NAC technologies.

Telephony

Telephony is the use of telephone technologies to connect computers. Historically, telephony referred to the technical elements related to analog voice-based telephone systems. However, it has grown and now includes the technical elements supporting both analog and digital voice systems, along with fax and data transmissions. Although plain old telephone service (POTS) transmission wires still exist in many areas of the world, much of the underlying infrastructure supports broadband digital transmissions.

An extension of telephony is Internet telephony, also known as Voice over Internet Protocol (VoIP). VoIP phones allow users to make phone calls using a network connection with access to the Internet, rather than traditional phone systems.

Dial-Up RAS

Dial-up Remote Access Service (RAS) uses telephony technologies, including phones and modems. Both the client and server need access to phone lines, and each must have a modem. Dial-up RAS allows the client to have access to a remote network over traditional phone wires.

As a simple example, I was a traveling trainer for many years. While I was on the road, I was able to dial in to the RAS server using a laptop with a modem anytime I needed access to the company network. Once I connected, I could access resources on the network similar to how I could access the resources if I was at my desk at work—just not as quickly because I was using a 56K dial-up modem.

The client has a modem and access to a phone line and can dial directly into the RAS server. The RAS server also has a modem and access to a phone line. Once connected, the RAS server provides access to the internal network.

The primary protocol used for dial-up access is Point-to-Point Protocol (PPP). When it was developed, tapping phone lines was considered rare, so PPP didn't include much security. However, PPP is often combined with other protocols to enhance security of the connection.

Long-distance phone costs can make dial-up cost prohibitive. VPNs provide better security than a dial-up solution and can reduce phone costs.

VPNs and VPN Concentrators

A *virtual private network (VPN)* allows a connection to a private network over a public network. The public network is most commonly the Internet, but it can also be a semiprivate leased line from a telecommunications company. Because the telecommunications company will often lease access to one physical line to several companies, the leased line is not truly private.

Access over a public network is a core security concern with VPNs. Different tunneling protocols encapsulate and encrypt the traffic to protect the data from unauthorized disclosure.

In large organizations, the VPN server is a VPN concentrator. A VPN concentrator includes all the services needed by a VPN server, including strong encryption and authentication techniques, and supports a large number of clients.

Connecting via Remote Access

Figure 4.8 shows three examples of how users can connect to internal networks from remote locations. The first VPN client connects to the Internet using a broadband connection to an Internet Service Provider (ISP), and the second VPN client connects to the Internet using a dial-up connection to an ISP. The third client bypasses the Internet and instead uses a dial-up connection directly to the RAS server.




Figure 4.8: Connecting to a VPN server and a RAS server

The figure shows the VPN server and RAS server between two firewalls (configured as a DMZ). The firewalls provide some protection for the servers and the private network. The VPN server is reachable through a public IP address, making it accessible from any other host on the Internet and the RAS server is reachable through a phone line. Both the VPN server and RAS server use an internal RADIUS server to authenticate clients before granting them access to the internal network.

Of course, the firewall access control lists (ACLs) need to include rules allowing the traffic. The ports you open are dependent on the tunneling protocols used for the VPN.

IPsec as a Tunneling Protocol

Chapter 3 introduced Internet Protocol security (IPsec) as a method of encrypting traffic on the wire. IPsec supports Tunnel mode, which encrypts the entire IP packets, and is the mode used with VPNs. Transport mode only encrypts the payload within the IP packets and is used within private networks.

IPsec provides security in two ways:

- **Authentication.** IPsec includes an Authentication Header (AH) to allow each of the hosts in the IPsec conversation to authenticate with each other before exchanging data. AH provides authentication and integrity. AH uses protocol ID number 51.
- **Encryption.** IPsec includes Encapsulating Security Payload (ESP) to encrypt the data and provide confidentiality. ESP includes AH so it provides authentication, integrity, and confidentiality. ESP uses protocol ID number 50.

The term *protocol ID number* may look like a typo, but it isn't. AH and ESP are identified with protocol ID numbers, not port numbers. Chapter 3 presented routers and firewalls. You may remember from Chapter 3 that a basic packet-filtering firewall can filter packets based on IP addresses, ports, and some protocols, such as Internet Control Message Protocol (ICMP) and IPsec. Packet filters use the protocol ID numbers to identify AH and ESP traffic.

IPsec uses Internet Key Exchange (IKE) over port 500 to authenticate clients in the IPsec conversation. IKE creates security associations (SAs) for the VPN and uses these to set up a secure channel between the client and the VPN server.

Remember this

IPsec is a secure encryption protocol used with VPNs. Encapsulating Security Payload (ESP) provides confidentiality, integrity, and authentication for VPN traffic. IPsec uses Tunnel mode for VPN traffic and can be identified with protocol ID 50 for ESP. It uses IKE over port 500.

L2TP and IPsec

Cisco and Microsoft joined forces to create the Layer 2 Tunneling Protocol (L2TP). Although it is possible to use L2TP to create a tunnel between devices, L2TP doesn't include any encryption, so it does not provide confidentiality of the data. However, you can combine IPsec with L2TP (as L2TP/IPsec) to provide security for the VPN tunnel. L2TP uses UDP port 1701.

NAT and IPsec

IPsec and Network Address Translation (NAT) are not compatible with each other. NAT manipulates the IP header of the packets when it translates the IP addresses. This change causes the receiving end of the VPN tunnel to discard the packet as invalid.

If the path to the VPN server is through a device using NAT, you need to look for alternatives. NAT Traversal (NAT-T) is one possible choice, or you could use another tunneling protocol.

TLS and SSL

Some tunneling protocols use either Secure Sockets Layer (SSL) or Transport Layer Security (TLS) to secure the VPN channel. As an example, Secure Socket Tunneling Protocol (SSTP) encrypts VPN traffic using SSL over port 443. Using port 443 provides a lot of flexibility for many administrators and rarely requires opening additional firewall ports. It is a useful alternative when the VPN tunnel must go through a device using NAT, and IPsec is not feasible. OpenVPN and OpenConnect are two open source applications that can use TLS to create a secure channel.

PPTP

Many Microsoft implementations used Point-to-Point Tunneling Protocol (PPTP) for VPNs. PPTP uses Microsoft's Point-to-Point Encryption to create the secure channel, but PPTP is rarely used today because of known vulnerabilities. PPTP uses TCP port 1723.

Remember this

L2TP is a VPN tunneling protocol and it is commonly combined with IPsec (as L2TP/IPsec). L2TP uses UDP port 1701. PPTP uses TCP port 1723.

Site-to-Site VPNs

A site-to-site VPN includes two VPN servers that act as gateways for two networks separated geographically. For example, an organization can have two locations. One is its headquarters and the other is a remote office. It can use two VPN servers to act as gateways to connect the networks at the two locations together.

A benefit of the site-to-site model is that it connects both networks without requiring additional steps on the part of the user. Users in the remote office can connect to servers in the headquarters location as easily as if the servers were in the remote office. Connecting to the remote server may be slower than connecting to a local server, but otherwise it's transparent to end users.

In contrast, in a host-to-gateway model, the end user makes the direct connection to the VPN server and is very much aware of the process.

VPN Over Open Wireless

Public wireless hot spots typically have little or no security. Their goal is to provide users with free wireless access and often include a warning indicating that the network doesn't have any security. Attackers with wireless sniffers can typically capture any traffic sent over the network.

The best protection for users, especially when entering confidential data, such as usernames and passwords, is to ensure data is encrypted. One method is to ensure that connections are using Hypertext Transfer Protocol Secure (HTTPS) connections.

Another method is to use a VPN service that provides secure VPN connections over open wireless connections. For example, Private Internet Access and TunnelBear are two applications that provide this service, and applications are available for most platforms, including smartphones.

Network Access Control

Allowing access to your private network can expose your network to a significant number of risks from the clients. If an employee VPNs into the network with a malware-infected computer, this computer can then infect other computers on the internal network. Network access control (NAC) methods provide continuous security monitoring by inspecting computers and preventing them from accessing the network if they don't pass the inspection.

Most administrators have complete control over computers in their network. For example, they can ensure the clients have up-to-date antivirus software installed, operating systems have current patches applied, and their firewalls are enabled. However, administrators don't have complete control of computers employees use at home or on the road.

NAC provides a measure of control for these other computers. It ensures that clients meet predetermined characteristics prior to accessing a network. NAC systems often use *health* as a metaphor, indicating that a client meets these predetermined characteristics. Just as doctors can quarantine patients with certain illnesses, NAC can quarantine or isolate unhealthy clients that don't meet the predefined NAC conditions.

Inspection and Control

Administrators set predefined conditions for healthy clients and those that meet these preset conditions can access the network. The NAC system isolates computers that don't meet the conditions. Common health conditions checked by a NAC are:

- Up-to-date antivirus software, including updated signature definitions
- Up-to-date operating system, including current patches and fixes
- Firewall enabled on the client

NAC clients have authentication agents (sometimes called health agents) installed on them. These agents are applications or services that periodically check different conditions on the computer and document the status in a statement of health. When a client connects to a NAC-controlled network, the NAC system queries the client's authentication agent and this agent responds with a statement of health.

Consider Figure 4.9. When a VPN client accesses the network, the VPN server queries the NAC health server to determine required health conditions. The VPN server also queries the client for a statement of the client's health. As long as the client meets all health requirements, NAC allows the client to access the network.




Figure 4.9: Using network access control

However, if a client doesn't meet the health conditions mandated by the NAC server, the VPN server redirects the client to a remediation network (also called a quarantine network). The remediation network includes resources the client can use to get healthy. For example, it would include current approved patches, antivirus software, and updated virus signatures. The client can use these resources to improve its health and then try to access the network again.

While NAC can inspect the health of VPN clients, you can also use it to inspect the health of internal clients. For example, internal computers may occasionally miss patches and be vulnerable. NAC will detect the unpatched system and quarantine it. If you use this feature, it's important that the detection is accurate. In at least one situation, the NAC identified healthy clients as unhealthy and prevented these healthy systems from accessing the network.

Similarly, your organization may allow visitors or employees to plug in their mobile computers to live wall jacks for connectivity, or connect to a wireless network. NAC inspects the clients, and if they don't meet health conditions, they may be granted Internet access through the network but remain isolated from any other network activity.

Remember this

Network access control (NAC) includes methods (such as health agents) to inspect clients for health, such as having up-to-date antivirus software. NAC can restrict access of unhealthy clients to a remediation network. You can use NAC for VPN clients and for internal clients.

Chapter 4 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Understanding IDSs and IPSs

- Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) inspect traffic using the same functionality as a protocol analyzer.
- A host-based IDS (HIDS) can detect attacks on local systems such as workstations and servers. The HIDS protects local resources on the host and can detect some malware that isn't detected by traditional antivirus software.
- A network-based IDS (NIDS) detects attacks on networks.
- A signature-based IDS uses signatures to detect known attacks or vulnerabilities. Vendors create a number to identify each signature, or use the number in the Common Vulnerabilities and Exposures (CVE) list.
- An anomaly-based (also called heuristic-based or behavior-based) IDS requires a baseline and detects attacks based on anomalies or when traffic is outside expected boundaries.
- A false positive sends an alert indicating an attack when an attack is not active. False positives increase the workload of administrators. A false negative is when an attack is active, but not reported.
- Honeypots and honeynets appear to have valuable data and attempt to divert attackers away from live networks. Security personnel use them to observe current attack methodologies and gather intelligence on attacks.
- An intrusion prevention system (IPS) is similar to an active IDS except that it's placed in-line with the traffic, and can stop attacks before they reach the internal network. An IPS can actively monitor data streams, detect malicious content, and mitigate the effect of malicious activity.
- IDSs and IPSs can also protect internal private networks, such as private supervisory control and data acquisition (SCADA) networks.

Securing Wireless Networks

- You can limit the coverage of a wireless access point (WAP) to a single room or a building by reducing the power level or modifying the placement of the antenna. This can help prevent unauthorized users from connecting to the wireless network. You can increase the wireless footprint by increasing power levels.
- Most WAPs have omnidirectional antennas. A Yagi antenna is a high-gain directional antenna and you can connect two buildings with WAPs and Yagi antennas.
- Site surveys identify the footprint of a wireless network and potential threats. Administrators perform site surveys during the planning stage, and perform periodic site surveys to identify threats.
- Wired Equivalent Privacy (WEP) is an older wireless protocol that is not secure. Wi-Fi Protected Access (WPA) was an initial improvement over WEP and it uses Temporal Key Integrity Protocol (TKIP) with Rivest Cipher 4 (RC4), which is compatible with older hardware.
- WEP implemented RC4 incorrectly using a small initialization vector (IV). IV attacks often use packet injection to generate traffic and crack the encryption key.
- WPA cracking attacks capture the four-way authentication handshake and then perform a brute force attack to discover the passphrase.
- Wi-Fi Protected Access II (WPA2) is the current standard and it supports Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP). CCMP is based on the strong Advanced Encryption Standard (AES) encryption protocol.
- WPA/WPA2 Personal mode uses a preshared key (PSK). It is easy to implement and is used in many smaller wireless networks.
- WPA/WPA2 Enterprise mode is more secure than Personal mode because it adds authentication. It uses an 802.1x authentication server implemented as a RADIUS server.
- 802.1x servers use one of the Extensible Authentication Protocol (EAP) versions, such as Protected EAP (PEAP), EAP-Tunneled Transport Layer Security (EAP-TTLS), EAP-TLS, or Lightweight EAP (LEAP).
- The most secure EAP method is EAP-TLS, and it requires a certificate on the server and on each of the wireless clients. PEAP and EAP-TTLS require a certificate on the server, but not the client. PEAP is often implemented with Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2). LEAP is proprietary to Cisco and does not require a certificate.
- WAPs in hot spots often use Isolation mode to segment, or separate, wireless users from each other. Additionally, they disable security to make it easy for users to connect or use WEP with Open System Authentication.
- You can restrict access to wireless networks with media access control (MAC) filtering. However, an attacker with a wireless sniffer can discover authorized MACs and perform a spoofing attack.
- A wireless audit checks WAP power levels, antenna placement, wireless footprint, and encryption techniques. It often includes war driving techniques and can detect rogue access points.
- Disabling the service set identifier (SSID) broadcast (the network name) hides a wireless network from casual users. However, an attacker with a wireless sniffer can easily determine the SSID even if SSID broadcasting is disabled.

- A rogue access point is an unauthorized WAP. Attackers can use them to capture data on your network. Unauthorized users can access your network via a rogue access point.
- An evil twin is a rogue access point using the same SSID as an authorized WAP.
- Bluejacking involves sending unsolicited messages to a phone. Bluesnarfing involves accessing data on a phone such as email and contact lists.

Exploring Remote Access

- VPNs provide remote access to an internal network for mobile users. Firewall ACLs include rules to allow VPN traffic based on the tunneling protocol.
- VPN concentrators provide secure remote access to a large number of remote users.
- Some VPNs use IPsec to encrypt data in a secure tunnel during transit.
- IPsec is a common tunneling protocol used with VPNs. It can secure traffic in a site-to-site tunnel and from clients to the VPN. IPsec uses Tunnel mode for VPNs. ESP encrypts VPN traffic and provides confidentiality, integrity, and authentication.
- Firewalls identify IPsec ESP traffic with protocol ID 50 and AH traffic with protocol ID 51. IKE creates the security association for the IPsec tunnel and uses port 500.
- Other tunneling protocols include SSTP (using SSL over port 443), L2TP (over UDP port 1701), and PPTP (over TCP port 1723).
- NAC inspects clients for specific health conditions such as up-to-date antivirus software, and can redirect access to a remediation network for unhealthy clients. NAC can be used with VPN clients and with internal clients.

Chapter 4 Practice Questions

1. Which of the following network tools includes sniffing capabilities?
 - A. IDS
 - B. WAP
 - C. VPN
 - D. NAC
2. A HIDS reported a vulnerability on a system using an assigned vulnerability identification number. After researching the number on the vendor's web site, you identify the recommended solution and begin applying it. What type of HIDS is in use?
 - A. Network-based
 - B. Signature-based
 - C. Heuristic-based
 - D. Anomaly-based
3. Management is concerned about malicious activity on your network and wants to implement a security control that will detect unusual traffic on the network. Which of the following is the BEST choice to meet this goal?
 - A. Network firewall
 - B. Signature-based IDS
 - C. Anomaly-based IDS
 - D. Honeypot
4. Administrators have noticed an increased workload recently. Which of the following can cause an increased workload from incorrect reporting?
 - A. False negatives
 - B. False positives
 - C. Separation of duties
 - D. Signature-based IDSs
5. A security company wants to identify and learn about current and new attack methodologies. Which of the following is the BEST choice to meet this objective?
 - A. Pen test
 - B. HIDS
 - C. Honeypot

D. Firewall logs

6. Of the following choices, what can you use to divert malicious attacks on your network away from valuable data to worthless fabricated data?

- A. IPS
- B. Proxy server
- C. Web application firewall
- D. Honeypot

7. Your network IDS recently detected an attack on a server. Upon investigation, you discover that the IDS does not have a signature on this attack. Instead, the IDS detected it using a heuristic analysis. Of the following choices, what is the MOST likely category of this attack?

- A. Definition
- B. CVE
- C. Zero-day
- D. Phishing

8. You need to provide connectivity between two buildings without running any cables. You decide to use two WAPs and a high-gain directional antenna. Which of the following antennas is the BEST choice to meet this need?

- A. Yagi
- B. Omni
- C. Isotropic
- D. Dipole

9. You are assisting a user implement a wireless network in his home. The wireless hardware he has requires the RC4 protocol. What type of security is BEST for this network?

- A. WEP
- B. WPA-TKIP
- C. WPA-AES
- D. WPA2 Enterprise

10. You want to implement the STRONGEST level of security on a wireless network. Which of the following supports this goal?

- A. Implementing WEP
- B. Disabling SSID broadcast
- C. Enabling MAC filtering
- D. Implementing WPA2

11. You are planning to deploy a WLAN and you want to ensure it is secure. Which of the following provides the BEST security?

- A. WEP Enterprise
- B. WPA2 TKIP
- C. SSID broadcast
- D. WPA2 CCMP

12. Your organization is planning to implement a wireless network using WPA2 Enterprise. Of the following choices, what is required?

- A. An authentication server with a digital certificate installed on the authentication server
- B. An authentication server with DHCP installed on the authentication server
- C. An authentication server with DNS installed on the authentication server
- D. An authentication server with WEP running on the access point

13. You are assisting a small business owner in setting up a public wireless hot spot for her customers. Which of the following actions are MOST appropriate for this hot spot?

- A. Enabling Open System Authentication
- B. Enabling MAC filtering
- C. Disabling SSID broadcast

D. Installing Yagi antennas

14. Homer is able to connect to his company's wireless network with his smartphone but not with his laptop computer. Which of the following is the MOST likely reason for this disparity?

- A. His company's network has a MAC address filter in place.
- B. His company's network has enabled SSID broadcast.
- C. His company's network has enabled CCMP.
- D. His company's network has enabled WPA2 Enterprise.

15. Management asks you if you can modify the wireless network to prevent users from easily discovering it. Which of the following would you modify to meet this goal?

- A. CCMP
- B. WPA2 Enterprise
- C. SSID broadcast
- D. MAC address filter

16. A war driver is capturing traffic from a wireless network. When an authorized client connects, the attacker is able to implement a brute force attack to discover the encryption key. What type of attack did this war driver use?

- A. WPS attack
- B. IV attack
- C. Packet injection
- D. WPA cracking

17. Your organization hosts three wireless networks for different purposes. A recent site survey audit discovered the information shown in the following table:

SSID	Security	Channel	Power
GetCertifiedVisitors	WPA2	1	71 dBm
GetCertifiedEmployee	WPA2	2	94 dBm
GetCertifiedEmployees	WPA2	3	73 dBm
GetCertifiedKiosk	WPA2	5	79 dBm

What does this indicate?

- A. Evil twin
- B. Rogue access point
- C. Interference
- D. Near field communication

18. An attacker is able to access email contact lists on your smartphone. What type of attack is this?

- A. Bluesnarfing
- B. War chalking
- C. War driving
- D. Bluejacking

19. Your organization is planning to implement a VPN and wants to ensure it is secure. Which of the following protocols is the BEST choice to use with the VPN?

- A. HTTP
- B. SFTP
- C. IPsec
- D. PPTP

20. An automated process isolated a computer in a restricted VLAN because the process noticed the computer's antivirus definitions were not up to date. What is the name of this process?

- A. NFC

- B. NIPS
- C. NIDS
- D. NAC

Chapter 4 Practice Question Answers

1. **A.** Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) include sniffing capabilities allowing them to inspect packet streams for malicious activity. None of the other tools have the capability of inspecting packets. A wireless access point (WAP) provides access to a wired network for wireless devices. A virtual private network (VPN) provides access to an internal network for remote users. A network access control (NAC) system inspects clients to ensure they meet minimum security requirements.
2. **B.** If the issue has an assigned number, it must be known, so it is signature-based. A host-based intrusion detection system (HIDS) is not network-based. A heuristic-based (or anomaly-based) detection system catches issues that are not previously known.
3. **C.** An anomaly-based (also called heuristic or behavior-based) detection system compares current activity with a previously created baseline to detect any anomalies or changes. A network firewall blocks and allows traffic, but does not detect unusual traffic. Signature-based systems use signatures similar to antivirus software. A honeypot is a server designed to look valuable to an attacker and can divert attacks.
4. **B.** False positives can cause an increased workload because they falsely indicate an alert has occurred. A false negative doesn't report an actual attack, so it doesn't increase the workload because administrators are unaware of the attack. Separation of duties ensures a single person can't control an entire process, so it is unrelated to increased workload. Signature-based intrusion detection systems (IDSs) don't necessarily cause an increased workload unless they have a high incidence of false positives.
5. **C.** A honeypot is a server designed to look valuable to an attacker and can help administrators learn about zero-day exploits, or previously unknown attacks. Security personnel perform a pen test (or penetration test) to determine if attackers can exploit existing vulnerabilities, but attackers may not try to do so. A host-based intrusion detection system (HIDS) attempts to detect intrusions on an individual host, but may not catch new methods against the network. Firewall logs can log connections, but don't identify new attack methods.
6. **D.** A honeypot can divert malicious attacks to a harmless area of your network, such as away from production servers holding valid data. An intrusion prevention system (IPS) can block attacks, but it doesn't divert it. A proxy server can filter and cache content from web pages, but doesn't divert attacks. A web application firewall (WAF) is an additional firewall designed to protect a web application.
7. **C.** Heuristic analysis has the best chance of detecting a zero-day attack. A zero-day attack is one that is unknown to vendors and because this attack doesn't have a signature, it is most likely unknown. Definition-based intrusion detection systems (IDSs) are the same as signature-based IDSs. Many signatures are based on the Common Vulnerabilities and Exposures (CVE) list. A phishing attack is an email, not an attack on a server.
8. **A.** A Yagi antenna is a high-gain directional antenna with a very narrow radiation pattern and is an ideal choice for this scenario. An isotropic antenna is theoretical and indicates the signal goes in all directions equally. Omnidirectional and dipole antennas attempt to mimic an isotropic antenna, but have stronger gains horizontally than vertically, assuming they are standing vertically.
9. **B.** Temporal Key Integrity Protocol (TKIP) uses RC4 and is compatible with older hardware so Wi-Fi Protected Access (WPA) with TKIP is the best option for this network. Wired Equivalent Privacy (WEP) uses RC4, but it is not secure and should not be used. WPA with Advanced Encryption Standard (AES) is stronger, but it uses AES instead of RC4. Wi-Fi Protected Access II (WPA2) Enterprise requires an 802.1x server and does not use RC4.
10. **D.** Wi-Fi Protected Access II (WPA2) provides the strongest level of security of the available answers. Wired Equivalent Privacy (WEP) is weak and should not be used. Disabling service set identifier (SSID) broadcast hides the network from casual users, but attackers can still discover it because the SSID is still included in some packets in plaintext. Attackers can bypass media access control (MAC) address filtering by spoofing authorized MAC addresses.
11. **D.** Wi-Fi Protected Access II (WPA2) with Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP) provides the best security of those listed. Wired Equivalent Privacy (WEP) is not secure and is not available in Enterprise mode. CCMP is stronger than Temporal Key Integrity Protocol (TKIP). Service set identifier (SSID) broadcast indicates the network name is broadcast, but this doesn't provide any security. If SSID broadcast is disabled, it hides the network from casual users, but attackers can still see it.
12. **A.** WPA2 Enterprise requires an 802.1x authentication server and most implementations require a digital certificate installed on the server. The network will likely have Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) services, but it isn't necessary to install them on the authentication server. Wired Equivalent Privacy (WEP) provides poor security and is not compatible with WPA2 Enterprise.
13. **A.** Open System Authentication is the best choice of those given for a public wireless hot spot. It is used with Wired Equivalent Privacy (WEP), doesn't require users to enter a preshared key or passphrase, and doesn't require the business owner to give out this information. It's also possible to disable security for the hot spot. Media access control (MAC) address filtering would be very difficult to maintain. Disabling service set identifier (SSID) broadcasting would make it difficult to find the wireless network, and installing a directional Yagi antenna isn't appropriate for a hot spot that

needs an omnidirectional antenna.

14. **A.** A media access control (MAC) address filter allows (or blocks) devices based on their MAC addresses, so it is likely that the filter is allowing Homer's smartphone but not allowing his laptop computer. Enabling the service set identifier (SSID) makes the network easier to see by casual users, but it does not block access even if SSID broadcast is disabled. Wi-Fi Protected Access II (WPA2) and Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP) both provide strong security, but they do not differentiate between devices.

15. **C.** You can disable service set identifier (SSID) broadcasting to prevent users from easily discovering the wireless networks. None of the other methods hide the network. Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP) provides stronger security for Wi-Fi Protected Access II (WPA2) and WPA2 Enterprise adds authentication for a wireless network. Media access control (MAC) address filtering can restrict access to the wireless network.

16. **D.** A Wi-Fi Protected Access (WPA) cracking attack captures traffic and then performs an offline brute force attack to discover the encryption key. Wi-Fi Protected Setup (WPS) attacks also use a brute force attack, but do not need to wait for an authorized client to connect. Initialization vector (IV) attacks often use packet injection techniques to generate more traffic in Wired Equivalent Privacy (WEP) attacks.

17. **B.** This indicates a rogue access point because the organization is hosting three wireless networks, but the survey found four. A rogue access point typically has a similar name (such as GetCertifiedEmployee in this example). An evil twin will have the exact name as an authorized WAP. An interference or jamming attack would make it difficult to connect to the access points causing users to disconnect often. Near field communication (NFC) refers to two devices communicating when they are close to each other and is unrelated to this scenario.

18. **A.** Attackers are able to access data (including email contact lists) on a smartphone in a bluesnarfing attack. War chalking is the practice of marking the location of wireless networks. War driving is the practice of looking for wireless networks, often by driving around. Bluejacking is the practice of sending unsolicited messages to other Bluetooth devices.

19. **C.** Internet Protocol secure (IPsec) is one of several protocols used to secure virtual private network (VPN) traffic. It is the best choice of the available answers. Hypertext Transfer Protocol (HTTP) doesn't provide any security. Secure File Transfer Protocol (SFTP) secures FTP transmissions but not VPNs. Point-to-Point Tunneling Protocol (PPTP) is an older protocol used with VPNs, but it is not as secure as IPsec.

20. **D.** Network access control is a group of technologies that can inspect systems and control their access to a network. In this scenario, NAC changed the computer's IP address to quarantine it in a restricted virtual local area network (VLAN). Near field communication (NFC) refers to standards that allow mobile devices to communicate with each other and is not related to VLANs. Network-based intrusion prevention systems (NIPSs) and network-based intrusion detection systems (NIDSs) protect a network from intrusions, but do not quarantine internal systems.

Chapter 5

Securing Hosts and Data

CompTIA Security+ objectives covered in this chapter:

- 1.3 Explain network design elements and components.**
 - Virtualization
 - Cloud computing (Platform as a Service, Software as a Service, Infrastructure as a Service, Private, Public, Hybrid, Community)
 - 1.4 Given a scenario, implement common protocols and services.**
 - Protocols (iSCSI, Fibre Channel, FCoE)
 - 2.1 Explain the importance of risk related concepts.**
 - Risks associated with Cloud Computing and Virtualization
 - 2.3 Given a scenario, implement appropriate risk mitigation strategies.**
 - Enforce policies and procedures to prevent data loss or theft
 - Enforce technology controls (Data Loss Prevention (DLP))
 - 2.6 Explain the importance of security related awareness and training.**
 - User habits (Personally owned devices)
 - 2.9 Given a scenario, select the appropriate control to meet the goals of security.**
 - Availability (Patching)
 - 3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques.**
 - Hardening (Disabling unnecessary services, Protecting management interfaces and applications, Password protection, Disabling unnecessary accounts)
 - Security posture (Initial baseline configuration)
 - 4.1 Explain the importance of application security controls and techniques.**
 - Application configuration baseline (proper settings)
 - Application hardening, Application patch management
 - 4.2 Summarize mobile security concepts and technologies.**
 - Device security (Full device encryption, Remote wiping, Lockout, Screen-locks, GPS, Application control, Storage segmentation, Asset tracking, Inventory control, Mobile device management, Device access control, Removable storage, Disabling unused features)
 - Application security (Key management, Credential management, Authentication, Geo-tagging, Encryption, Application whitelisting, Transitive trust/authentication)
 - BYOD concerns (Data ownership, Support ownership, Patch management, Antivirus management, Forensics, Privacy, On-boarding/off-boarding, Adherence to corporate policies, User acceptance, Architecture/infrastructure considerations, Legal concerns, Acceptable use policy, On-board camera/video)
 - 4.3 Given a scenario, select the appropriate solution to establish host security.**
 - Operating system security and settings, OS hardening, Patch management, White listing vs. black listing applications, Trusted OS, Host software baselining
 - Virtualization (Snapshots, Patch compatibility, Host availability/elasticity, Security control testing, Sandboxing)
 - 4.4 Implement the appropriate controls to ensure data security.**
 - Cloud storage, SAN, Handling Big Data
 - Data encryption (Full disk, Database, Individual files, Removable media, Mobile devices)
 - Hardware based encryption devices (TPM, HSM, USB encryption, Hard drive)
 - Data in-transit, Data at-rest, Data in-use, Permissions/ACL
 - 4.5 Compare and contrast alternative methods to mitigate security risks in static environments.**
 - Environments (SCADA, Embedded (Printer, Smart TV, HVAC control), Android, iOS, Mainframe, Game consoles, In-vehicle computing systems)
 - Methods (Network segmentation, Security layers, Application firewalls, Manual updates, Firmware version control, Wrappers, Control redundancy and diversity)
 - 5.2 Given a scenario, select the appropriate authentication, authorization or access control.**
 - Authentication (Trusted OS)
- **

In this chapter, you'll learn about different methods used to secure hosts, such as servers, workstations, and mobile devices. Steps include disabling unnecessary services and keeping them up to date. Data is one of the most valuable resources of any organization, and organizations take steps to protect data, prevent data leakage, and prevent the loss of confidentiality.

Implementing Host Security

Hosts include servers, workstations, and mobile computing devices. In an ideal world, they start in a secure state. Unfortunately, it's not an ideal world, and administrators need to be proactive to secure hosts before deployment and keep them secure after deployment. This section outlines several steps used to secure hosts, starting with basic hardening methods.

OS and Application Hardening

Hardening is the practice of making an operating system (OS) or application more secure from its default installation. An important first step when hardening operating systems and applications is to read the directions. The vendor documentation and guidelines include important details on steps you can take to secure them, but they're only useful when someone reads them and applies the knowledge.

Many vendors approach application development with a goal of usability over security. The application will be easy to set up and configure, but it may have gaping security holes. Other vendors value security over usability. It's secure when you install it, but it may not work as desired.

Interestingly, it's only when applications value security over usability that administrators are forced to read the documentation. If security settings prevent the application from working as desired, they have to dig into the documentation to modify the settings. In contrast, if they install it and it works out of the box, they may be called away to fight another crisis before checking the documentation for security issues.

Disabling Unnecessary Services

A core principle associated with hardening a system includes disabling or removing all unnecessary services. If a service is not running on a system, attackers cannot attack it, and it reduces the overall attack surface of the system. For example, an expert on exploiting File Transfer Protocol (FTP) vulnerabilities will be unsuccessful using these techniques on a server that is not running the FTP service. It doesn't matter how vulnerable a service is. If it's not running, attackers cannot exploit any of its vulnerabilities.

When you disable a service, you often remove access to the associated protocol. For example, if you disable the FTP service, you disable the FTP protocol. Some protocols, such as Transmission Control Protocol (TCP) and User Datagram Protocol (UDP), are necessary for connectivity within a network and cannot be disabled. Other protocols, such as FTP, Remote Desktop Protocol (RDP), and Simple Mail Transfer Protocol (SMTP), are optional application protocols supporting related services. If the server needs them, you'd enable them. For example, you'd enable SMTP on an email server, but would disable it on a server not sending or receiving email.

Disabling unnecessary services and removing unneeded protocols provide several key benefits, including the following:

- **Improves the overall security posture of systems.** Systems that are running only required services and protocols are less susceptible to attacks, which improves their overall security posture.
- **Reduces the attack surface.** Disabled services and protocols are not vulnerable to attacks, so by disabling unnecessary services and protocols, you reduce the system's attack surface. Both known attack methods and previously unknown zero-day vulnerability attack methods are unsuccessful when the services and protocols are disabled or removed. This includes known attack methods and previously unknown zero-day vulnerability attacks.
- **Reduces risks associated with open ports.** If an attacker does a port scan, the port scan fails on the associated port. For example, disabling the RDP service on a server causes a port scan on port 3389 to fail on the server, even if the port is open on a firewall between the attacker and the server.

Remember this

Hardening a server makes it more secure from its default installation. Disabling unnecessary services and protocols reduces the attack surface of a system and improves its overall security posture.

Eliminating Unneeded Applications

In addition to disabling unnecessary services to reduce vulnerabilities, it's important to uninstall unneeded software. Software frequently has bugs and vulnerabilities. Although patching software frequently closes these vulnerabilities, you can eliminate these vulnerabilities by simply eliminating unneeded applications.

Years ago, I was working at a small training company. One of the servers had a default installation of Windows. We were using the server as a file server, but because it wasn't hardened from the default installation, it was also running Internet Information Services (IIS), the Microsoft web server.

At some point, attackers released the Nimda virus, which exploited a vulnerability with IIS. Microsoft released a patch for IIS, but because IIS was installed by default and we weren't using it, we also weren't managing it. Ultimately, the Nimda virus found our server, and the worm component of Nimda quickly infected our network. If the IIS software hadn't been installed, the server would not have been vulnerable to the attack.

Disabling Unnecessary Accounts

Many operating systems and applications come with default accounts. A basic principle in hardening systems and applications is to disable unnecessary accounts before deploying the system. For example, the Guest account is disabled by default in current Windows systems, but in the past, administrators disabled it before deploying Windows systems.

Similarly, database server applications include default database accounts and have used blank or default passwords in the past. If the deployed application doesn't need the default accounts, administrators disable them before deployment. If the application needs the accounts, administrators

ensure the accounts have strong passwords.

Some applications also include backdoor accounts. A backdoor is an access point to an application or service that bypasses normal security mechanisms. Developers use backdoors for legitimate purposes to view the internal workings of an application or for ease of administration. However, the use of backdoors is strongly discouraged in the final released version. If a backdoor exists, you can expect attackers to locate and exploit it.

Protecting Management Interfaces and Applications

Many devices and applications have built-in tools used for administration. These tools often provide administrators with the ability to manage the devices, but their access needs to be protected to prevent unauthorized personnel from modifying them. One of the primary ways you can protect them is by disabling unnecessary accounts and changing the default passwords, as mentioned in the previous section.

As an example, Chapter 4, “Securing Your Network,” covers wireless networks and many wireless devices come with predefined accounts and passwords used in the application. A common combination is an Administrator account named “admin” and a password of “admin.” Unless this password is changed, an attacker can easily break into the network and cause considerable damage.

It’s also important to use best practices when creating and modifying passwords. For example, Chapter 1, “Mastering Security Basics,” includes information on creating strong passwords.

Using Baselines

A *baseline* is a known starting point and organizations commonly use baselines to provide known starting points for workstations and servers. One of the primary benefits of baselines is that they improve the overall security posture of systems. This works in three steps:

1. **Initial baseline configuration.** Administrators use various tools to deploy systems consistently in a secure state.
2. **Continuous security monitoring.** Automated tools monitor the systems for any baseline changes. Some tools such as vulnerability scanners monitor the systems and report any changes they detect. Other tools such as Group Policy automatically reconfigure the systems to the baseline settings when they detect changes.
3. **Remediation.** Chapter 4 covers network access control (NAC). NAC methods can detect some changes to baseline settings and automatically isolate or quarantine systems in a remediation network. Typically, administrators need to correct the problems in these systems manually.

There are several different types of baselines, including:

- Security baselines used to ensure systems start in a secure state
- Configuration baselines used to ensure systems are deployed consistently
- Host software baselines used to document software installed on host systems
- Application configuration baselines used to document appropriate application settings
- Performance baselines used to identify when system performance is degrading

Security Baselines

A *security baseline* is a secure starting point for an operating system or application. Organizations often identify the requirements for security baselines in a written security policy. Administrators then use the security policy as a guide when creating the baseline. For example, the security policy might require all desktop systems to have up-to-date antivirus software installed, host-based firewalls enabled, and FTP disabled. Administrators use different methods, such as Group Policy and imaging, to deploy the baseline. Later, they can check existing systems against the security baseline to verify the system is still secure.

Enforcing Security Baselines with Group Policy

Microsoft domains use Group Policy to standardize the configuration of systems. An administrator can create and apply a Group Policy Object (GPO) to configure all the systems in the domain, or target specific systems. Some common security settings applied by Group Policy include:

- **Account settings.** Administrators can configure several specific security settings such as disabling the Guest account and renaming the Administrator account.
- **Password and account lockout policies.** These settings ensure users maintain strong passwords and lock accounts to prevent an attacker from trying to guess the password.
- **Audit policies.** When enabled, auditing logs certain events, such as when users log on or off, or when they access objects such as files.
- **User rights.** Rights refer to what users can do on a system, such as log on using Remote Desktop or shut down the system. Administrators can configure these settings to allow or restrict these rights depending on the needs of the organization.
- **System services.** These settings allow administrators to disable services such as FTP.
- **Software restrictions.** Administrators use these settings to control what software can be installed on a system and what software can run on a system. For example, they can use these settings to prevent the installation or use of peer-to-peer (P2P) software.

The magic of Group Policy is that an administrator can configure a single setting within a GPO and apply it to multiple users or computers with very little effort. A GPO works the same way whether it's being applied to five systems or five thousand. Group Policy is applied when a computer starts up and when a user logs on. The system periodically checks to see if any Group Policy settings are changed and automatically applies these new settings.

Another benefit of Group Policy is that it regularly reapplies security settings. If a problem or attack compromises a system, this process helps keep the Group Policy security settings in place.

The first step in creating the security baseline is creating a written security policy. Once the organization creates the security policy, administrators use different methods, such as Group Policy, security templates, or imaging, to deploy the baseline. Later, they can check existing systems against the security baseline to verify the system is still secure.

For example, imagine that your organization's security policy mandates that users should not be able to install software. Administrators deploy systems enforcing this policy. Later, they can check existing systems to ensure that users cannot install software and the original security baseline is still intact.

An organization will typically have several security baselines. For example, end-user operating systems use one baseline, generic servers use

another baseline, and specialty servers use other baselines.

Each operating system is different so you won't find a standard checklist to lock down all operating systems. However, there is a place to check—the vendor's documentation. If you're trying to secure an operating system or an application running on the operating system, check the documentation. This documentation often includes valuable information with easy-to-follow steps.

Some vendors include tools to help create a security baseline. For example, Microsoft Server operating systems include the Security Configuration Wizard (SCW). SCW leads administrators through a series of questions about a system and then creates an Extensible Markup Language (XML) database file that includes a wide assortment of security settings. Administrators can import these settings into a Group Policy Object to apply them.

Using Imaging for Baselines

One of the most common methods of deploying systems is with images. An *image* is a snapshot of a single system that administrators deploy to multiple other systems. Imaging has become an important practice for many organizations because it streamlines deployments while also ensuring they are deployed in a secure manner. Figure 5.1 and the following text identify the overall process of capturing and deploying an image:




Figure 5.1: Capturing and deploying images

1. Administrators start with a blank source system. They install and configure the operating system, install and configure any desired applications, and modify security settings. Administrators perform extensive testing to ensure the system works as desired and that it is secure before going to the next step.
2. Next, administrators capture the image. Symantec Ghost is a popular imaging application, and Windows Server 2012 includes free tools many organizations use to capture and deploy images. The captured image is simply a file that can be stored on a server or copied to external media, such as a DVD or external USB drive.
3. In step 3, administrators deploy the image to multiple systems. When used within a network, administrators can deploy the same image to dozens of systems during an initial deployment, or to just a single system to rebuild it. The image installs the same configuration on the target systems as the original source system created in step 1.

Administrators will often take a significant amount of time to configure and test the source system. They follow the same hardening practices discussed earlier and often use security and configuration baselines. If they're deploying the image to just a few systems such as in a classroom setting, they may create the image in just a few hours. However, if they're deploying it to thousands of systems within an organization, they may take weeks or months to create and test the image. Once they've created the image, they can deploy it relatively quickly with very little administrative effort.

Imaging provides two important benefits:

- **Secure starting point.** The image includes mandated security configurations for the system. Personnel who deploy the system don't need to remember or follow extensive checklists to ensure that new systems are set up with all the detailed configuration and security settings. The deployed image retains all the settings of the original image. Administrators will still configure some settings, such as the computer name, after deploying the image.
- **Reduced costs.** Deploying imaged systems reduces the overall maintenance costs and improves reliability. Support personnel don't need to learn several different end-user system environments to assist end users. Instead, they learn just one. When troubleshooting, support personnel spend their time focused on helping the end user rather than trying to learn the system configuration. Managers understand this as reducing the total cost of ownership (TCO) for systems.

Many virtualization tools include the ability to convert an image to a virtual system. In other words, once you create the image, you can deploy it to either a physical system or a virtual system. From a security perspective, there is no difference how you deploy it. If you've locked down the image for deployment to a physical system, you've locked it down for deployment to a virtual system.

Imaging isn't limited to only desktop computers. You can image any system, including servers. For example, consider an organization that maintains 50 database servers in a large data center. The organization can use imaging to deploy new servers or as part of its disaster recovery plan to restore failed servers. It is much quicker to deploy an image to rebuild a failed server than it is to rebuild a server from scratch. As long as administrators keep

the images up to date, this also helps ensure the recovered server starts in a secure state.

Remember this

Standardized images include mandatory security configurations. This ensures systems start in a secure state and reduces overall costs. Administrators are able to identify anomalies by comparing settings, services, and applications in the image with settings, services, and applications on live computers.

Configuration Baselines

A *configuration baseline* identifies the configuration settings for a system. This includes settings such as printer configuration, application settings, and TCP/IP settings. This is especially useful when verifying proper operation of a system. As an example, if a server is no longer operating correctly, it might be due to a configuration change. Administrators might be able to identify the problem by comparing the current settings against the baseline and correcting any discrepancies.

The differences between a configuration baseline and a security baseline can be a little fuzzy. The security baseline settings are strictly security related. The configuration baseline settings ensure consistent operation of the system. However, because the configuration baseline contributes to improved availability of a system, which is part of the security triad, it also contributes to overall security.

An important consideration with a configuration baseline is keeping it up to date. Administrators should update the configuration baseline after changing or modifying the system. This includes after installing new software, deploying service packs, or modifying any other system configuration settings.

U.S. Government Configuration Baseline (USGCB)

The U.S. government has been using standard images for many years. This started as a Standard Desktop Core Configuration (SDCC) with the U.S. Air Force and morphed into the Federal Desktop Core Configuration (FDCC) mandated by the Office of Management and Budget (OMB) for all federal agencies. The current version is the United States Government Configuration Baseline (USGCB), which includes images for Windows 7 and Linux. Before using these images, many agencies were repeating common security errors. Smaller agencies without extensive information technology (IT) or security experience deployed systems without locking them down. Even when government security professionals knew about common attacks and how to protect systems, the smaller agencies didn't have the expertise or manpower to implement the fixes.

However, agencies are now consistently deploying new systems in a secure state. The images include the mandated security settings and do not require extensive security knowledge or expertise to deploy. Additionally, these images include settings that are compliant with Security Content Automation Protocol (SCAP). SCAP validation functionality is built in to many vulnerability scanners, allowing them to verify the security settings haven't been modified.

Of course, these aren't the only security measures the agencies take, but they do provide a good start.

...

Host Software Baselines

A *host software baseline* lists all software installed on a system, along with a list of approved software. Administrators sometimes refer to it as an application baseline. Administrators can perform automated scans of computers to create an inventory of all the installed applications. They can compare this with a list of approved applications to identify unauthorized applications. This type of scan can also identify applications that aren't up to date with current patches and are vulnerable to attack.

Remember this

Host software baselines provide a list of approved software and a list of software installed on systems. Administrators can use this to identify unauthorized software installed on systems. Unauthorized software is not maintained and can easily become vulnerable without patching.

Application Configuration Baselines

Application configuration baselines identify the proper settings for applications. They are similar to system configuration baselines with the exception that they only refer to settings for specific applications. As an example, imagine an organization hosts several database servers running Microsoft SQL Server. Administrators configure the settings on these systems in a secure state and document the settings. Later, administrators can audit the servers to ensure the settings haven't been modified.

Performance Baselines

A *performance baseline* identifies the overall performance of a system at a point in time. If performance deteriorates later, administrators can compare the current performance against the baseline report. The differences between the current measurements and the baseline help an administrator differentiate between normal performance and actual problems.

The baseline report includes information on usage of basic system hardware resources, such as the processor, memory, disk, and network interface card (NIC). It also includes additional system data, such as logs to show normal behavior.

As an example, Performance Monitor is a tool used within Windows systems to create performance baseline reports. A performance baseline report captures snapshots of key metrics every 30 minutes throughout a seven-day period. These snapshots give a good picture of a system's performance during peak performance times and slack times. An administrator can later compare current performance with the baseline to identify any differences.

Baseline Reporting

Baseline reporting is the process of comparing systems against a baseline to identify discrepancies or anomalies. It can be used with any type of baseline. As an example, a security baseline configures systems in a known secure state. Later, administrators can audit the systems to ensure they are still in the same known secure state.

Several tools are available to assist with baseline reporting. Vulnerability scanners scan systems looking for specific security settings and provide a report for systems with different security settings. Similarly, application baseline reporting includes a scan of systems to identify installed applications and compares it with authorized applications. If it finds unauthorized software, such as P2P software, the baseline report lists the software along with computer's name and other relevant details on the computer such as its IP address.

Remember this

Baseline reporting provides a report after comparing baselines with current systems. Administrators use baseline reporting for security baselines, operating system baselines, application configuration baselines, and software baselines.

Whitelisting Versus Blacklisting Applications

Whitelisting and blacklisting are two additional methods used to protect hosts, including workstations, servers, and mobile devices. A *whitelist* is a list of applications authorized to run on a system. A *blacklist* is a list of applications the system blocks.

You can use Software Restriction Policies in Microsoft Group Policy for both whitelisting and blacklisting for computers within a domain. For a whitelist, you identify the applications that can run on the system, and Group Policy blocks all other applications. For a blacklist, you identify the applications that cannot run on the system, and Group Policy allows any other applications. For example, if users have been running a specific type of unauthorized P2P software, or operating system games such as FreeCell, you can add these applications to the blacklist. Group Policy will then prevent them from running.

Some antivirus software supports the use of whitelists. For example, Kaspersky Lab maintains a whitelist database. This list helps prevent false positives where antivirus software incorrectly identifies valid applications as malicious software. Note that in this example, the whitelist doesn't include all the safe software in the world. Instead, it includes a list of applications that are known to be safe. Antivirus software doesn't need to check these applications as closely as unknown applications.

Remember this

Application whitelisting identifies authorized software for workstations, servers, and mobile devices. It prevents users from installing or running software that isn't on the list.

Trusted OS

A *trusted operating system* (trusted OS) meets a set of predetermined requirements with a heavy emphasis on authentication and authorization. The overall goal of a trusted operating system is to ensure that only authorized personnel can access data based on their permissions. Additionally, a trusted operating system prevents any modifications or movement of data by unauthorized entities. With this in mind, a trusted OS helps prevent malicious software (malware) infections because it prevents malicious or suspicious code from executing.

A trusted OS meets a high level of security requirements imposed by a third party. For example, the Common Criteria for Information Technology Security Evaluation (or simply Common Criteria) includes requirements for a trusted OS. Operating systems that meet these requirements can be certified as trusted operating systems.

Understanding Virtualization

Virtualization is a technology that has been gaining a lot of popularity in recent years. It allows you to host one or more virtual systems, or virtual machines (VMs), on a single physical system. With today's technologies, you can actually host an entire virtual network within a single physical system and organizations are increasingly using virtualization to reduce costs.

When discussing VMs and studying for the CompTIA Security+ exam, you should understand the following terms:

- **Hypervisor.** The software that creates, runs, and manages the VMs is the hypervisor. Several virtualization technologies currently exist, including VMware, Microsoft Hyper-V, Windows Virtual PC (VPC), and Oracle VM VirtualBox. All of these have their own hypervisor software.
- **Host.** The physical server hosting the VMs is the host. It requires more resources than a typical system, such as multiple processors, massive amounts of RAM, fast and abundant hard drive space, and one or more fast network cards. Although these additional resources increase the cost of the host, it is still less expensive than paying for multiple physical systems. It also requires less electricity, less cooling, and less physical space.
- **Guest.** Operating systems running on the host system are guests or guest machines. Most hypervisors support several different operating systems, including various Microsoft operating systems and various Linux distributions. Additionally, most hypervisors support both 32-bit and 64-bit operating systems.
- **Patch compatibility.** It's important to keep VMs patched and up to date. Patches applied to physical systems are compatible with virtual systems.
- **Host availability/elasticity.** Elasticity refers to the ability to resize computing capacity based on the load. For example, imagine one VM has increased traffic. You can increase the amount of processing power and memory used by this server relatively easily. This allows you to ensure it remains available even with the increased demand.

Snapshots

Snapshots provide you with a copy of the VM at a moment in time, which you can use as a backup. If the VM develops a problem, you can revert the image to the state it was in when you took the snapshot. You are still able to use the VM just as you normally would. However, after taking a snapshot, the hypervisor keeps a record of all changes to the VM.

Administrators commonly take snapshots of systems prior to performing any risky operation. Risky operations include applying patches or updates, and installing new applications. Ideally, these operations do not cause any problems, but occasionally they do. By creating snapshots before these operations, administrators can easily revert the system to the previous state.

Sandboxing and Security Control Testing

A *sandbox* is an isolated area and is often used for testing programs. When using VMs, you can create them so that they are isolated in a sandbox environment. Sandboxing prevents the VMs from interacting with any other VMs, the physical host, or devices on the network. The term comes from a sandbox in a playground. Children can play in the sandbox where they are relatively safe (and parents can easily keep their eyes on them). Similarly, application developers can play in a virtual sandbox knowing that any changes they make will not affect anything outside the sandbox.

Administrators and security professionals also use sandboxing to test various security controls before deploying them to a live production network. Virtualization provides a high level of flexibility when testing security controls because the environments are easy to re-create. For example, they can test the effectiveness of antivirus software to detect malware released within a sandbox. If the antivirus software doesn't detect the malware and the malware causes problems, it is easy to revert the system to a previous state using a snapshot. Also, the isolation within the sandbox prevents the malware from spreading.

Similarly, virtualized sandboxes are useful for testing patches. For example, software vendors typically develop software updates and patches, but they need to test them in various environments before releasing them. They could create VMs for multiple operating systems. When they're ready to test, they turn on one of the VMs, take a snapshot, and then apply and test the patch. If the patch causes a problem, they can easily revert the VM.

Remember this

Virtualization allows multiple virtual servers to operate on a single physical server. It provides increased availability with lower operating costs. Additionally, virtualization provides a high level of flexibility when testing security controls, updates, and patches because they can easily be reverted using snapshots.

VMs as Files

It's worth pointing out that virtual machines are simply files. These files certainly have some complexity, but still, they are just files. As an example, files associated with a Hyper-V VM include:

- **VHD files.** These files hold data on the virtual hard disk (VHD) within the VM. For example, if the VM has a C and a D drive, these files contain the contents of the C and D drives.
- **XML files.** These files contain the configuration details of the VM within Extensible Markup Language (XML) files. Each VM includes an associated XML file and each snapshot includes an associated XML file identifying the configuration at that time.
- **AVHD files.** These are automatic VHDs, more commonly known as differencing disks. They hold the differences between the current disk and a previously created snapshot.
- **VSV files.** These files hold the saved state for devices associated with the VM. Instead of shutting down a VM logically, it's also possible to save the current state and turn it off. This is similar to using Hibernate mode in a laptop. Hibernate mode saves the current state of the laptop in a file and turns the laptop off. When you power it back on, it returns to the exact same state as it was before going into Hibernate mode.
- **BIN files.** These files hold the memory for systems that are in a saved state.

Because the VM is just a group of files, it becomes relatively easy to move them from one physical server to another. For example, if one of your physical servers becomes overloaded, you can move virtual servers off the overloaded system to another physical server. Some virtual server management software makes this as simple as dragging and dropping the virtual servers from one host to another.

It's also easy to restore a failed virtual server. If you create a backup of the virtual server files and the original server fails, you simply restore the files. You can measure the amount of time it takes to restore a virtual server in minutes. In contrast, rebuilding a physical server can take hours.

Many virtualization products allow you to manage multiple virtual systems on a single server, even when the virtual servers are running on separate physical hosts. For example, you might have five physical servers hosting three virtual servers each, and you can manage all of them through a single management interface. This includes taking snapshots, reverting snapshots, and moving the virtual servers from one physical host to another.

Networking Connectivity

Organizations usually configure online virtual servers so that they can communicate with other virtual and physical systems on the network. They use virtual network interface cards (NICs), virtual switches, and virtual networks for connectivity. These are all contained within the physical host.

Both Microsoft and VMware support the use of virtual local area networks (VLANs) with virtual switches. Just as you can use VLANs to segment traffic on a physical network, you can also use VLANs to segment traffic on a virtual network.

However, it's also possible to configure the virtual systems so that they are completely isolated. For example, you can isolate a virtual server so that it can't communicate with any other virtual or physical systems. In this way, it works just like a single system without a NIC. You can also group several virtual servers in their own virtual network so that they can communicate with each other but are isolated from hosts on the physical network.

Many security professionals use virtual systems and virtual networks to test and investigate malware. Malware released in an isolated environment presents minimal risk to the hardware and host operating system. Unfortunately, some malware is able to detect that it is running in a virtual environment. In some cases, malware developers have written code to change the behavior of the malware when it discovers it is running in a virtual environment.

Remember this

Virtual local area networks (VLANs) separate or segment traffic on physical networks. You can also create VLANs using virtual switches within a virtual environment hosted on a physical server.

Risks Associated with Virtualization

Despite the strengths of virtualization technologies, you should understand some weaknesses. Many people consider virtual machine escape (VM escape) to be the most serious threat to virtual system security. Loss of confidentiality and loss of availability can also be a concern.

VM Escape

VM escape is an attack that allows an attacker to access the host system from within the virtual system. The host system runs an application or process called a hypervisor to manage the virtual systems. In some situations, the attacker can run code on the virtual system and interact with the hypervisor.

Most virtual systems run on a physical server with elevated privileges, similar to administrator privileges. A successful VM escape attack often gives the attacker unlimited control over the host system and each virtual system within the host.

When vendors discover VM escape vulnerabilities, they write and release patches. Just as with any patches, it is important to test and install these patches as soon as possible. This includes keeping both the physical and the virtual servers patched.

Loss of Confidentiality

As a reminder, each virtual system or virtual machine is just one or more files. Although this makes it easy to manage and move virtual machines, it also makes them easy to steal.

It's worth pointing out that a virtual machine includes the operating system and data, just as a physical system would have both the operating system

and data on its physical drives. For example, a virtual machine can include a database with credit card data, company financial records, or any type of proprietary data.

With this in mind, consider an administrator who has turned to the dark side and become a malicious insider. The insider has access to the systems and can easily copy the virtual machine, take it home, and launch it on another physical server. At this point, the attacker has access to the system and the data.

You may remember from Chapter 1 that one of the primary methods of protecting against loss of confidentiality is with encryption. Virtual systems support encryption just as physical systems do. If any of the data is important, you can protect it with encryption.

Implementing Patch Management

Software is not secure. There. I said it. As someone who has written a few programs over the years, that's not easy to say. In a perfect world, extensive testing would discover all the bugs, exploits, and vulnerabilities that cause so many problems.

However, because operating systems and applications include millions of lines of code, testing simply doesn't find all the problems. Instead, most companies make a best effort to test software before releasing it. Later, as problems crop up, companies write and release patches or updates. Administrators must apply these patches to keep their systems up to date and protected against known vulnerabilities.

Patch management ensures that systems and applications stay up to date with current patches. This is one of the most efficient ways to reduce operating system and application vulnerabilities because it protects systems from known vulnerabilities. Patch management includes a group of methodologies and includes the process of identifying, downloading, testing, deploying, and verifying patches.

As an example, many web browsers have bugs that make them vulnerable to drive-by downloads. When successful, a web site will download and install malicious software after a user does nothing more than visit a specially crafted web page. This happens without the user's knowledge or approval. Attackers trick users into visiting these pages using a variety of methods, including sending links in spam emails. One click and the user's system is infected. In contrast, if administrators ensure web browsers and operating systems are up to date, these bugs are patched and the drive-by downloads fail.

Remember this

Patch management procedures ensure that operating systems and applications are up to date with current patches. This protects systems against known vulnerabilities.

Automated Versus Controlled Deployment

Individual users and small organizations sometimes configure systems to automatically download and install patches. These systems periodically check for updates and when they're available, they automatically download them. Then they typically install the patches in the middle of the night.

Unfortunately, some patches cause problems with systems. For example, a patch to Windows 7 in August 2014 caused some systems to crash into the infamous blue screen of death (BSOD). The systems automatically rebooted, crashed again, and continued an endless rebooting loop. This only happened to systems that had a specific OpenTypeFont installed and in a nonstandard font directory. However, for the people who met these specifics, the result was quite troubling. Imagine if your organization had 50 such computers. You'd walk in one day with 50 users screaming for help.

Because of the potential for problems such as this, larger organizations take control of deploying patches. Their patch management program plans for the release of patches and includes methods of testing and deploying them.

Scheduling Patch Management

Microsoft releases most of its patches on the second Tuesday of the month (known as Patch Tuesday). IT departments can plan on this release so that they can immediately begin evaluating the patches, testing relevant patches, and deploying them. Organizations with certain partnerships with Microsoft often receive advance notice of these patches. This allows them to plan for them before Patch Tuesday. Microsoft will sometimes release an out-of-band (OOB) patch that is released right away, but it only does this for critical vulnerabilities.

Other operating systems, such as Unix and Linux, don't currently release patches on a schedule, but they still release patches. Administrators can sign up for notifications about patches and plan their timeline based on these notifications.

Immediately after Microsoft releases patches on Patch Tuesday, many attackers go to work. They read as much as they can about the patches, download them, and analyze them. They often attempt to reverse engineer the patches to determine exactly what the patch is fixing.

Next, the attackers write their own code to exploit the vulnerability on unpatched systems. They often have exploits attacking systems the very next day—Exploit Wednesday. Because many organizations take more than a single day to test the patch before applying it, this gives the attackers time to attack unpatched systems. For organizations without a patch management program, it gives attackers much longer to attack unpatched systems.

Additionally, some attackers discover unknown exploits before Patch Tuesday. They recognize that Microsoft will be releasing patches on the second Tuesday of the month, so they wait until the second Wednesday before launching major attacks to exploit the vulnerability. Unless Microsoft releases an out-of-band patch, this gives them a full month to exploit systems before a patch is available.

Testing Patches

Patches can fix one problem but create others, such as an endless rebooting loop. Consider the worst-case scenario. In some unfortunate situations, systems shut down and never work again. If this happens to your home computer, it is inconvenient. However, if one thousand computers within an organization stop working one day, it can be catastrophic.

Organizations avoid this problem by testing patches before deploying them. The goal of testing is to ensure that a patch does not introduce new problems. For testing to be realistic, you need to install the patch on systems that mirror the production environment. In other words, if all the users have

new computers, it won't do any good to test a patch on an older system.

Regression testing is a specific type of testing used to detect any new errors (or regressions). In regression testing, administrators run a series of known tests on a system and compare the results with previously run tests.

Deploying and Verifying Patches

After testing the patches, administrators deploy them. They don't deploy the patches manually though. Instead, they use systems management tools to deploy the patches in a controlled manner. For example, Microsoft ConfigMgr is a systems management tool used for many purposes, including patch management.

In addition to deploying patches, systems management tools also include a verification component that verifies patch deployment. They periodically query the systems and retrieve a list of installed patches and updates. They then compare the retrieved list with the list of deployed patches and updates, providing reports for any discrepancies. In some networks, administrators combine this with network access control (NAC) technologies and isolate unpatched systems in quarantined networks until they are patched.

Mitigating Risk in Static Environments

Static computing environments are relatively constant, especially when compared with typical computers connected to a network. Historically, administrators didn't see a need to provide much protection to these environments, but as with just about anything security related, things change. The need to protect them has become clearer in recent years. Some examples of static environments are:

- **Supervisory control and data acquisition (SCADA) systems.** These are typically industrial control systems within large facilities such as power plants or water treatment facilities. These systems are typically contained within isolated networks that do not have access to the Internet.
- **Embedded systems.** This includes computing components embedded in printers, smart televisions, and heating, ventilation, and air conditioning (HVAC) control systems. Although these aren't typical attack vectors, they could be. For example, suppose an attacker was able to remotely turn off the HVAC system or trick it into keeping the temperature at 95 degrees within a data center. The resulting damage to systems within this data center could be catastrophic.
- **Mobile systems.** This includes devices running Android and Apple iOS operating systems. However, due to the growth in their use and capabilities, these are becoming less static.
- **Mainframes.** Mainframe computers are high-powered systems usually performing dedicated functions within an organization. Mainframes might be contained within isolated networks. However, it is more common for them to be connected to an organization's primary network so that personnel can access them.
- **Game consoles.** Gaming has become big business and has driven many of the advances in memory, graphics, and processor power. Today, gaming consoles have powerful processors for games, but also for connectivity with networks and other players, and players often have the ability to make in-game purchases through the systems. If the consoles hold credit card or other user data, they are susceptible to attacks just like any other computer.
- **In-vehicle computing systems.** Most of today's cars have powerful processors within them. Some ensure the engine is running smoothly and efficiently. Other embedded systems provide direct connectivity to services such as OnStar, satellite radio, and the Internet. In one reality show, police set a trap with a car in an area where several cars were stolen. A thief was in for a big surprise when he stole it, though. After breaking into the car, he was soon driving away. The police sent a signal to the car, turning its engine off, and quickly apprehended the thief. However, if the police can do this, what happens when criminals learn how to do it, too?

Understanding Stuxnet

A great example of the need to protect embedded systems comes from Stuxnet. Stuxnet is a computer worm designed to attack a specific embedded system, used in one of Iran's nuclear enrichment facilities. It caused centrifuges to spin fast enough to tear themselves apart and some reports indicated it destroyed as many 20 percent of these centrifuges.

Security expert Roel Schouwenberg completed extensive research on Stuxnet and identified how it operated in six major steps:

1. **Infection.** Stuxnet first infected Windows systems through infected USB drives after someone plugged one into the system. One of the architects of Stuxnet reportedly said "...there is always an idiot around who doesn't think much about the thumb drive in their hand." Indeed, USB sticks have been the source of many infections.
2. **Search.** Stuxnet checks the network of the infected system looking for the targeted system.
3. **Update.** If it finds the targeted system, it downloads an updated version of the worm.
4. **Compromise.** It then attempts to compromise the targeted system. When first released, Stuxnet took advantage of four zero-day vulnerabilities. Zero-day vulnerabilities are either unknown to the vendor, or the vendor hasn't released a patch for them yet.
5. **Control.** It then sends signals to the systems. A late version of Stuxnet told the systems to spin the centrifuges uncontrollably.
6. **Deceive and destroy.** While it was causing the centrifuges to spin out of control, it was sending false data to engineers monitoring the system. Monitoring systems indicated everything was fine.

Protecting Static Systems

There are several methods recommended to mitigate risks in static environments. They include:

- **Redundancy and diversity controls.** Redundancy controls ensure a system continues to operate even when it suffers a failure. For example, a redundant array of inexpensive disks (RAID) includes one or more extra or redundant disks that take over if another disk fails. SCADA systems often include redundant controls to take over if one fails. Diversity refers to protecting systems with diverse security controls. This is often done by using security controls from different vendors. For example, it's common to create a demilitarized zone (DMZ) with two firewalls, but use firewalls from different vendors, providing a diverse defense.
- **Network segmentation.** By placing systems in separate networks, it protects them from potentially malicious traffic in a primary network.

An extreme form of network segmentation is removing the systems from any access to the primary network. For example, SCADA systems can be connected to each other, but not to any other network. Chapter 3, “Understanding Basic Network Security,” covers virtual local area networks (VLANs) and VLANs provide another method of network segmentation.

- **Security layers.** Defense-in-depth methods ensure that systems have multiple layers of security. For example, installing firewalls to block unauthorized traffic into a SCADA network provides one layer of security. Installing a network intrusion prevention system (NIPS) provides an additional layer of security for a SCADA network. The NIPS can inspect data streams for malicious traffic and block it. Layered security helps protect a system even if an attacker is able to breach one layer of security.
- **Application firewalls.** Application firewalls can inspect traffic and identify specific commands within a protocol. Although they are sometimes difficult to implement in a full network supporting multiple protocols, they can be quite effective at protecting static environments that only support a minimum number of protocols.
- **Manual updates.** One challenge with automatic updates is that they require frequent access to the Internet. By using manual updates, it allows administrators to download the updates on a separate environment, and verify they are valid before applying the updates to systems in static environments.
- **Firmware version control.** Most static systems have embedded firmware installed on them. When vendors discover bugs or security flaws, they write and release firmware updates to correct the issue. Firmware version control is a management control that ensures systems are periodically examined to verify the firmware is up to date with the most current version.
- **Wrappers.** Some operating systems such as Linux and Unix use Transmission Control Protocol (TCP) wrappers to filter traffic coming in and out of a system. This is similar to how a host-based firewall uses an access control list (ACL) to filter traffic. It’s possible to use similar techniques in some embedded systems to filter traffic.

Remember this

Incorporating control redundancy and diversity into security designs is a key method of protecting static environments such as supervisory control and data acquisition (SCADA) systems. Networks holding SCADA systems can be protected using virtual local area networks (VLANs) to segment traffic and network-based intrusion protection systems (NIPS) to block unwanted traffic.

Securing Mobile Devices

Mobile devices are smartphones, tablets, and laptop computers. Because they are mobile, they are more susceptible to some threats such as theft, and data security is one of the primary concerns with mobile devices.

There are many methods available to reduce risks associated with mobile devices. Some methods focus on securing the device, while others focus on application security. The following sections cover these topics along with some concerns related to users bringing their own device into an organization.

Some general security concepts include:

- **Encryption.** Encryption protects against loss of confidentiality on multiple platforms, including workstations, servers, mobile devices, and data transmissions. Encryption methods such as full device encryption provide device security, application security, and data security.
- **Authentication and device access control.** Authentication methods such as usernames and passwords ensure only authorized personnel can access devices. Laptops support multifactor authentication, which is especially useful when the laptop includes valuable data. Smartphones and tablets typically have short passcodes, but users can be forced to use more secure authentication methods when they access the organization's network.
- **Device access control.** Authentication methods protect access to devices. Without proper authentication methods, attackers can bypass device access controls.

Device Security

The primary scenario that concerns mobile device users and administrators is loss or theft of a device. For example, if a thief steals Homer's smartphone or tablet, Homer will have two primary concerns. He'll want to locate the device and prevent the thief from accessing any data on the device. Several tools address both of these concerns.

The primary tool used to locate a lost device is Global Positioning System (GPS), which pinpoints its location. GPS capabilities are standard on smartphones and tablets. As an example, Apple products use iCloud features and can help you locate iPhones and iPads from any computer with Internet access. The Find My iPhone or Find My iPad feature needs to be enabled on the device before it is lost. You can then log onto the iCloud site with a web browser.

Figure 5.2 shows a screenshot of the Find My iPhone web page after locating an iPad. I expanded the map to show more states, but you can easily home in on the exact street location. With just a few clicks, you can play a sound on the iPad (to help you find it), put it into Lost mode, or erase it.



Figure 5.2: Find My iPhone

Lost mode locks it with a passcode and displays a message on the screen. If you enter your phone number, the message includes your phone number. If a Good Samaritan finds your device, she can use this phone number to call you and return your device. The Erase iPad selection sends a signal to erase all data and settings on the device. After everything is erased, it will display a custom message from you, along with your phone number. However, you won't be able to access any features to locate the device.

It's worthwhile mentioning that while GPS can be an effective tool to help you locate and recover a system, it can also be a vulnerability. An attacker may be able to use GPS to track the location of an individual who owns a mobile device. Because of this, many users disable GPS tracking.

The following list summarizes device security concepts and concerns:

- **Removable storage.** USB thumb drives and other removable storage devices are a source of data leakage and malware distribution so security policies often restrict the use of USB thumb drives and other portable devices such as music players. Disabling removable storage capabilities makes it more difficult for users to copy data to and from the devices and protects them from malware. When using external USB hard drives, encryption can be effective at protecting the confidentiality of the data. However, it's important to use strong access controls to ensure attackers cannot bypass the encryption and access the data.
- **Storage segmentation.** In some mobile devices, it's possible to segment storage of data. For example, users might be required to use external storage for any corporate data to reduce the risk of data loss if the device is lost or stolen.
- **Screen locks.** Most devices support the use of a passcode or password to lock the device. This is similar to a password-protected screen saver on desktop systems and prevents someone from easily accessing the device and the data it contains. You probably won't keep the thief out of your mobile device for very long, but you can slow someone down. This gives you time to send a remote wipe signal before the thief accesses the data. Some devices have an additional setting that erases all the data if the incorrect passcode is entered too many times.
- **Lockout.** Many devices include a lockout feature, such as iPad's Lost mode lock. It can send a signal to lock the device with a passcode, even if it didn't originally have a passcode.
- **Remote wiping.** This is similar to the Erase iPad feature. It sends a remote signal to the device to wipe or erase all the data. Remote wipe capabilities are useful if the phone is lost. The owner can send a remote wipe signal to the phone to delete all the data on the phone. This also deletes any cached data, such as cached online banking passwords, and provides a complete sanitization of the device by removing all valuable data.
- **Disabling unused features.** Basic hardening practices for desktop and server systems apply to mobile devices, too. If ports and protocols aren't needed, they should be disabled or removed. Similarly, if mobile devices have any features that are not needed, they should be disabled. A side benefit of disabling unused features is that it reduces the drain on the battery and helps the battery to last longer.
- **Asset tracking.** Mobile devices are easy to lose track of so organizations often use asset-tracking methods to reduce losses. For example, when a user is issued a mobile device, asset-tracking methods record it. Similarly, if the user leaves the company, asset-tracking methods ensure the user returns the device.
- **Inventory control.** Many organizations use automated methods for inventory control. For example, radio-frequency identification (RFID) methods can track the movement of devices. These are the same types of devices used in stores to prevent shoplifting. If someone exits without paying, the RFID device transmits when the shoplifter gets close to the exit door, and sounds an alarm. Organizations won't

necessarily have an alarm, but they can track the movement of devices.

Remember this

Mobile device security includes device encryption to protect the data, screen locks to help prevent unauthorized access, and remote wipe capabilities to delete all data on a lost phone. Radio-frequency identification (RFID) methods are often used for inventory control.

BYOD Concerns

Many organizations are adopting bring your own device (BYOD) policies. These policies allow users to connect their own smartphones and tablets to the organization's network. Data security is a primary concern when allowing employee-owned devices onto a network, so different rules and guidelines are implemented to protect data. Additionally, there are several concerns that complicate these policies.

As an example, imagine that Homer has a tablet that he uses for email, Internet research, and various other personal uses. If his company allows him to connect this tablet to the corporate network, it will make it easier for Homer to access company email and other company resources needed for his job. However, is Homer's tablet protected from threats? Does it have antivirus malware installed? Does Homer keep it up to date with current patches? Who owns the data on the tablet? Who provides support if he has problems connecting to the network? Most of these issues would be addressed in an acceptable use policy, or in some cases a separate BYOD policy. The following topics summarize the BYOD concerns:

- **Acceptable use policy.** Chapter 11, "Exploring Operational Security," covers an acceptable use policy (AUP) in more depth, but as an introduction, it describes user responsibilities when using an organization's IT resources. Employees are required to periodically read and acknowledge the contents of the policy. When used with a BYOD policy, the AUP will define the responsibilities of employees if they choose to connect their personal devices to the organization's network.
- **Adherence to corporate policies.** If users are unwilling to adhere to the corporate policies, they are not allowed to connect their devices to the network.
- **Privacy.** Users have a right to privacy, but not all activities are private. For example, it isn't appropriate for users to connect to the network for Internet access and then spend their time gambling. AUPs often include a privacy statement informing users what activities they can consider private.
- **User acceptance.** These policies can potentially cause problems with some users. The attitude is that the device is their own, so the organization doesn't have any business managing it. However, when users understand the risks associated with unmanaged devices, they are often more willing to accept the policies. If users are unwilling to accept the policies, the organization simply blocks them from connecting their device to the company network.
- **Data ownership.** Typically, the policy will specify that all data owned by the organization remains the organization's data. This normally includes email sent through the organization's network.
- **Support ownership.** Support can be tricky. From a user's perspective, it's only one device. However, from the IT department's perspective, it is a multitude of smartphones and tablets. If the IT department must support them all, it can create a monumental workload. Because of this, many organizations limit the types of devices supported in a BYOD policy.
- **Architecture/infrastructure considerations.** Instead of allowing users to have full connectivity with an organization's network, it is possible to segment these devices to a restricted access network. For example, many users want access to the Internet for email and web browsing, but don't necessarily need to connect to servers within the network. Organizations can set up VLANs that segment these users. The VLANs can have a path to the Internet with appropriate security, but block access to the primary network.

Remember this

Data security is a significant concern related to BYOD policies. You can use VLANs to isolate mobile devices from the primary network, while still granting them access to the Internet.

- **Forensics.** Forensic analysts have tools to analyze mobile devices just as they have tools to analyze traditional computers. For example, they can identify deleted files. However, a company security specialist doesn't necessarily have permission to perform a forensic analysis on an employee-owned device.
- **Legal concerns.** If an organization doesn't take the time to define BYOD, AUP, and privacy policies, it can result in legal issues. For example, if a user has been copying sensitive data onto an employee-owned device, the organization might face a lawsuit if it attempts to confiscate the device to remove the data. It isn't possible to avoid all legal issues with policies, but they go a long way to ensure people understand their responsibilities, and help avoid most legal issues.
- **On-boarding/off-boarding.** *On-boarding* refers to the procedures to allow users to connect their devices to the network, and *off-boarding* refers to the procedures that remove the devices from the network. The most important consideration is simply to have established procedures. For example, having employees read and acknowledge the BYOD policy might be all that is required for on-boarding. When employees leave the organization, administrators remove all access to the network, including access with their mobile devices.

- **On-board camera/video.** Mobile devices have excellent cameras that can take both still pictures and videos. These are useful and valuable. However, it also allows personnel to secretly take pictures and videos. They might take pictures of other employees, screen images, or confidential data. Organizations that support BYOD policies need to consider if they want to restrict the use of the camera, or merely define restrictions in an AUP.

Mobile Device Management

Mobile device management (MDM) includes the technologies to manage mobile devices such as smartphones and tablets. The goal is to ensure these devices have similar security methods in place as desktop computers.

Traditional management tools such as Microsoft ConfigMgr ensure systems are kept up to date with current patches, have antivirus software installed with up-to-date definitions, and are secured using standard hardening practices. In the past, many of these management tools didn't include support for mobile devices. As BYOD became more and more popular, vendors upgraded these tools (and developed new ones) to manage mobile devices. ConfigMgr 2012 R2 includes support for many mobile devices, including Apple iOS-based devices and Android-based devices.

MDM tools often include the following features:

- **Patch management.** Patch management ensures that mobile devices are kept up to date with current patches.
- **Antivirus management.** Antivirus management ensures systems have antivirus software installed and it is up to date with current definitions.
- **Application control.** Some MDM tools can restrict what applications can run on mobile devices. They often use application whitelists to control the applications.

When employee-owned devices are in use, MDM tools typically block access to the network if the device doesn't meet minimum requirements. For example, if the device isn't patched or doesn't have up-to-date antivirus software, the MDM software works with network access control (NAC) technologies to prevent the device from connecting to the network.

Remember this

Mobile device management tools help ensure systems are up to date with current patches and have up-to-date antivirus installed. These tools often block devices that are not up to date.

Application Security

Application security can be a significant concern when it includes authentication methods. Many mobile devices store or cache credentials such as usernames and passwords, so if attackers can access the device, they might be able to access applications without providing credentials.

One credential management method is to prevent the use of cached or stored credentials. This requires users to reenter their credentials each time they log on. Similarly, key management methods prevent the use of cached or stored encryption keys.

These methods don't block or weaken authentication. Additionally, they don't prevent single sign-on methods using transitive trusts. They just force the user to enter cached credentials when they first access an application.

Geo-tagging adds geographical information to files such as pictures when posting them to social media web sites. For example, when you take a picture with a smartphone that has GPS features enabled, the picture application adds latitude and longitude coordinates to the picture. Thinking of friends and family, this is a neat feature. However, thinking of thieves and criminals, they can exploit this data. For example, if Lisa frequently posts pictures of friends and family at her house, these pictures identify her address. If she later starts posting pictures from a vacation location, thieves can realize she's gone and burglarize her home.

Remember this

Geo-tagging adds geographical information to files such as pictures when posting them on social media sites. Criminals can exploit this information when watching a specific person.

Protecting Data

Data is one of the most valuable resources any organization manages, second only to its people. If you ever tune into the news, you've likely heard about one of these stories. Unfortunately, data breaches are frequent and they affect millions of people. In the worst-case scenarios, thieves use the stolen data to empty bank accounts, rack up fraudulent charges on credit cards, and steal individuals' identities.

Some data breaches become big news. For example, in November and December of 2013, attackers hacked into Target's network and stole credit card data and personal information on more than 110 million customers. This attack was huge and media outlets reported on it for weeks. However, smaller data breaches occur almost daily.

The Identity Theft Resource Center tracks data breaches and regularly reports summaries. As of August 2014, they reported 480 data breaches exposing more than 17 million customer records during 2014. You can pick any month of any recent year and find examples of data breaches.

Losing control of data directly affects the reputation, and often the bottom line, of an organization. The importance of taking steps to protect valuable data cannot be overstated.

Chapter 11 covers security policies that an organization can implement to protect data. The security policy helps an organization classify and label its data. This section presents many of the security controls an organization can use to protect data based on the requirements set within a data security policy.

Comparing Data Categories

You will frequently see data categorized based on how it is used or stored. The most common terms are data at rest, data in transit, and data in use:

- *Data at rest* is any data stored on media. This includes data on hard drives, mobile phones, USB flash drives, external drives, and backups. Data can be stored as individual files or full databases. The best way to protect data at rest from an attacker is to encrypt it.
- *Data in transit* (or data in motion) is any data traveling over a network. Data loss prevention (DLP) techniques are effective at analyzing and detecting sensitive data sent over a network. You can also encrypt traffic sent over the network using encryption protocols, such as IPsec, SSH, or SFTP.
- *Data in use* refers to any data that resides in temporary memory. Applications retrieve stored data, process it, and may either save it back to storage or send it over a network. The application is responsible for protecting data in use.

Confidentiality is primarily protected through encryption and strong access controls. Chapter 1 focuses on access controls starting with strong authentication methods. This chapter discusses software-based and hardware-based encryption methods, and Chapter 10, “Understanding Cryptography,” covers specific encryption algorithms used to protect data.

Remember this

The primary methods of protecting the confidentiality of data (including data at rest and data in transit) are with encryption and strong access controls.

Protecting Confidentiality with Encryption

As mentioned in Chapter 1, one of the primary ways you can prevent the loss of confidentiality is by encrypting data. This includes encrypting data at rest no matter what type of device it is stored on and encrypting data in motion no matter what type of transmission media is used. It is much more difficult for an attacker to view encrypted data than it is to view data stored as plaintext.

You can use other tools to restrict access to data, but this isn't always effective. For example, consider the Microsoft NT File System (NTFS), which allows you to configure permissions within access control lists (ACLs). You can use NTFS to set permissions on files and folders to restrict access. However, if a thief steals a laptop with NTFS-protected files, it's a simple matter to access them. The thief simply moves the drive to another system as an extra drive, logs on as the administrator, and takes ownership of the files. Encryption isn't as easy to bypass.

Software-Based Encryption

Software-based encryption can encrypt individual files and folders, entire disks, removable media, mobile devices, and databases. Although software-based encryption is slower than hardware-based encryption, it is secure when using strong encryption algorithms.

File-Level Encryption

Many operating systems support file- and folder-level encryption. Linux systems support GNU privacy guard (gpg), which is a command-line tool used to encrypt and decrypt files with a password. Microsoft NTFS includes Encrypting File System (EFS), available in Windows Explorer. Users can right-click any file or folder, select Advanced, and select Encrypt Contents to Secure Data, as shown in Figure 5.3. An attacker will have a more difficult time accessing these encrypted files.

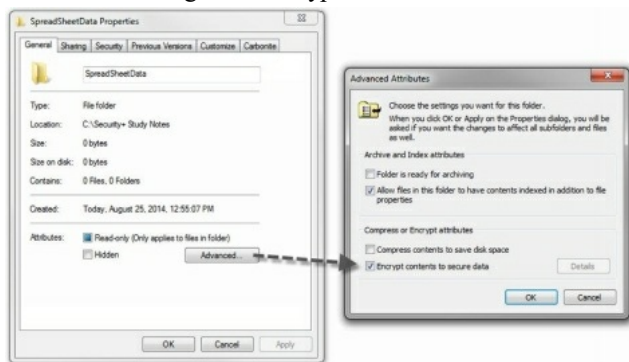


Figure 5.3: Encrypting a file with NTFS

A benefit of file- and folder-level encryption is that you can encrypt individual files without encrypting an entire disk. For example, a server may store files accessed by users throughout the company. Access controls provide a first level of protection for these files, but administrators may be able to bypass the access controls. Imagine that a company stores payroll data on the server and wants to ensure that a malicious insider with administrative privileges can't access the data. Using file encryption provides an additional level of protection.

One of the challenges with file-level encryption is that the encryption can be lost if an authorized user copies encrypted files to another disk that doesn't support encryption. For example, imagine that Bart encrypts a file on his system using NTFS and then copies the file to a FAT32-formatted USB drive. Because FAT32 doesn't support NTFS encryption, the system decrypts the file before copying it onto the drive. The solution to this is to use USBs with whole device or full disk encryption.

Full Disk Encryption

Full disk encryption programs encrypt an entire disk. For example, TrueCrypt is available on Linux and many other operating systems. It performs whole disk encryption for USB drives to protect the confidentiality of data if the device is lost. Users can access the data with a password, and TrueCrypt will decrypt and encrypt data on the fly without any other user intervention. It's also possible to use full disk encryption on traditional hard disk drives.

Encrypting Database Content

Another form of software-based encryption is with databases. For example, many database applications such as Oracle Database or Microsoft SQL Server include the ability to encrypt data held within a database. Although it's possible to encrypt the entire database, it's more common to encrypt specific data elements.

As an example, imagine a database includes a table named Customers. Each record within the table has multiple columns, including customer number, last name, first name, credit card number, and security code. Instead of encrypting the entire table, administrators can choose to encrypt only the credit card number and security code fields within each record. This protects the sensitive data, but doesn't waste valuable processing power encrypting data that isn't sensitive.

Remember this

File- and folder-level protection protects individual files. Full disk encryption protects entire disks, including USB flash drives and drives on mobile devices. Database column encryption protects individual fields within a database.

Hardware-Based Encryption

Although software-based encryption is useful, the drawback is that it can take extra processing power and time. It isn't as useful when a large quantity of data, such as an entire disk, needs to be encrypted or when performance is a concern.

You can use hardware-based encryption devices, such as a Trusted Platform Module or a hardware security module, for better performance. A significant benefit of hardware encryption is that it is much quicker than software encryption.

Table 5.1 provides an overview of these hardware encryption devices, and the following sections explore them in greater depth. Both use strong asymmetric encryption and provide a secure method of storing encryption keys.

Characteristics	TPM	HSM
Hardware	Chip in motherboard (included with many laptops)	Removable or external hardware device, (purchased separately)
Uses	Full disk encryption (for laptops and some servers)	High-end mission-critical servers (SSL accelerators, high availability clusters, certificate authorities)
Authentication	Performs platform authentication (verifies drive not moved)	Performs application authentication (only used by authorized applications)
Encryption Keys	Includes endorsement key (burned into chip) and storage root key Storage root key generates and protects other keys	Stores RSA keys used in asymmetric encryption and can generate keys

Table 5.1: A comparison of TPM and HSM features

Trusted Platform Module

A Trusted Platform Module (TPM) is a hardware chip on the computer's motherboard that stores cryptographic keys used for encryption. Many laptop computers include a TPM, but if the system doesn't include a TPM, it is not feasible to add one. Once enabled, the TPM provides full disk encryption capabilities. It keeps hard drives locked, or sealed, until the system completes a system verification and authentication process.

The TPM ships with a unique Rivest, Shamir, Adleman (RSA) key burned into it, which is used for asymmetric encryption. Additionally, the TPM can generate, store, and protect other keys used for encrypting and decrypting disks. TPMs use three categories of encryption keys:

- **Endorsement key.** The manufacturer embeds an endorsement key into the TPM. This stays with the TPM throughout its lifetime.
- **Storage root key.** The TPM creates the storage root key when a user adds a TPM owner password and activates the TPM. The TPM uses this key to create and protect other encryption keys used within applications.
- **Application keys.** These keys are derived from the storage root key and applications use them to encrypt disks. For example, Microsoft BitLocker uses an application key to encrypt entire disks.

If the system includes a TPM, you use an application within the operating system to enable it. For example, many Microsoft systems include BitLocker, which you can enable for systems that include the TPM.

BitLocker uses the TPM to detect tampering of any critical operating system files or processes as part of a platform verification process. Additionally, users provide authentication, such as with a smart card, a password, or a personal identification number (PIN). The drive remains locked until the platform verification and user authentication processes are complete.

If a thief steals the system, the drive remains locked and protected. An attacker wouldn't have authentication credentials, so he can't access the drive using a normal boot-up process. If the attacker tries to modify the operating system to bypass security controls, the TPM detects the tampering and keeps the drive locked. If a thief moves the drive to another system, the drive remains locked because the TPM isn't available.

Remember this

A Trusted Platform Module (TPM) is a hardware chip on the motherboard included on many newer laptops and it provides full disk encryption. A TPM includes a unique RSA asymmetric key. When a user activates the TPM, it creates a storage root key, which the TPM uses to generate and store other cryptographic keys.

Hardware Security Module

A hardware security module (HSM) is a security device you can add to a system to manage, generate, and securely store cryptographic keys. High-performance HSMs are external devices connected to a network using TCP/IP. Smaller HSMs come as expansion cards you install within a server, or as devices you plug into computer ports.

One of the noteworthy differences between an HSM and a TPM is that HSMs are removable or external devices. In comparison, a TPM is a chip embedded into the motherboard. You can easily add an HSM to a system or a network, but if a system didn't ship with a TPM, it's not feasible to add one later. Both HSMs and TPMs provide secure encryption capabilities by storing and using RSA keys. Many high-performance servers use HSMs to

store and protect keys.

Remember this

A hardware security module (HSM) is a removable or external device that can generate, store, and manage RSA keys used in asymmetric encryption. Many server-based applications use an HSM to protect keys.

Data Leakage

Another danger with data comes from data leakage. Just as a leak in one of your car's tires can let air out, leaving you with a flat tire, data can leak out of a company, leaving the company with flat financial performance. Worse, the data leakage can result in losses, reversing the company's profits.

Data exfiltration is the unauthorized transfer of data outside an organization and is a significant concern with data leakage. In some cases, attackers take control of systems and transfer data outside an organization using malware. It's also possible for malicious insiders to transfer data.

Data Loss Prevention

Data loss prevention (DLP) techniques examine and inspect data looking for unauthorized data transmissions. You may also see this term as data leak prevention. A DLP system can be network-based to inspect data in motion, storage-based to inspect data at rest, or endpoint-based to inspect data in use. In some scenarios, the DLP control prevents the use of hardware to prevent losses.

Data in Motion

Chapter 3 discusses different types of content filters used in unified threat management (UTM) devices, such as web security gateways. These devices monitor incoming data streams looking for malicious code. In contrast, a network-based DLP monitors outgoing data looking for sensitive data, specified by an administrator.

DLPs will scan the text of all emails and the content of any attached files, including documents, spreadsheets, presentations, and databases. Even if a user compresses a file as a Zip file before sending it, the DLP examines the contents by simply unzipping it.

As an example, I know of one organization that routinely scans all outgoing emails looking for Personally Identifiable Information (PII), such as Social Security numbers. The network-based DLP includes a mask to identify Social Security numbers as a string of numbers in the following format: ####-##-####. If an email or an attachment includes this string of numbers, the DLP detects it, blocks the email, and sends an alert to a security administrator.

Many organizations classify and label data using terms such as *Classified*, *Confidential*, *Private*, and *Sensitive*. It is easy to include these search terms in the DLP application, or any other terms considered important by the organization. Network-based DLPs are not limited to scanning only email. Many can scan the content of other traffic, such as FTP and HTTP traffic.

Endpoint Protection

Another method of preventing data loss is by restricting use of hardware at the computer (endpoint). This includes prohibiting the use of portable devices such as USB flash drives or preventing certain content from being printed.

Portable storage devices refer to any storage system that you can attach to a computer and easily copy data. It primarily refers to USB hard drives and USB flash drives, but many personal music devices, such as MP3 players, use the same type of flash drive memory as a USB flash drive. Users can plug them into a system and easily copy data to and from a system. Additionally, many of today's smartphones include storage capabilities using the same type of memory.

As mentioned several times throughout this book, USB drives represent significant risks to an organization. They can transport malware without the user's knowledge and can be a source of data leakage. Malicious users can copy and steal a significant amount of information using an easily concealable thumb drive. Users can misplace these drives, and the data can easily fall into the wrong hands.

Because of the risks, it's common for an organization to include security policy statements to prohibit the use of USB flash drives and other mobile storage devices. Some technical policies block use of USB drives completely.

A DLP solution is more selective and it can prevent a user from copying or printing files with specific content. For example, it's possible to configure a DLP solution to prevent users from copying or printing any classified documents marked with a label of *Confidential*. The DLP software scans all documents sent to the printer, and if it contains the label, the DLP software blocks it from reaching the printer.

Remember this

A network-based data loss prevention (DLP) system can examine and analyze network traffic. It can detect if confidential company data or any PII data is included in email and reduce the risk of internal users emailing sensitive data outside the organization. Similarly, endpoint DLP solutions can prevent users from copying or printing sensitive data.

Understanding SANs

Some organizations outgrow their storage needs and find that disk drives on a single server are no longer adequate. One possible upgrade is to a storage area network (SAN). A storage area network (SAN) provides connectivity for high-speed data storage devices in a specialized high-speed network. Other servers in the network connect to the SAN to access data. Because a SAN is dedicated to serving and storing data, it is generally much more efficient than a typical file server is.

Although hard drives are the most common, a SAN is able to host data on a variety of different media. For example, a SAN might include data stored on disks, tape, and optical media. However, the SAN typically includes capabilities beyond basic disks, CDs, and tapes.

It is common for SAN disks to be configured in arrays for better performance and fault tolerance. Tape libraries include the ability to automatically load and unload tapes via a robotic device. Similarly, optical jukeboxes include the ability to automatically load and unload optical media using a robotic device. Combined, all of this allows a SAN to store and serve a large amount of data.

In addition to optimized data storage, a SAN also provides high-speed transfers between devices within the SAN. Systems that connect to the SAN via the SAN network can also benefit from these high-speed transfers.

Virtual SANs (VSANs) are a newer feature that several vendors offer. Two goals are to lower overall storage costs associated with traditional SANs and to eliminate time lags or latencies associated with traditional networked storage.

Fibre Channel

Many SANs use Fibre Channel (FC) instead of a traditional TCP/IP network to provide high-speed transfers between the devices. Some FC networks enjoy transfer speeds of up to 16 gigabits per second.

Originally, FC only used fiber cable, but eventually support for copper cabling was added. Developers changed the name from the American English *fiber* to the British English *fibre* to indicate it supports both fiber and copper connections.

A drawback of FC networks is that they require specialized hardware and cabling. This often increases the costs, making an FC SAN too expensive for most organizations.

iSCSI

Internet Small Computer System Interface (iSCSI) transfers traditional SCSI commands over IP networks. A huge benefit is that it can utilize an existing network infrastructure to connect systems, without requiring specialized SCSI cabling. In other words, it's possible to implement an iSCSI SAN within local area networks (LANs) without the additional cost of specialized connectivity hardware.

FCoE

A Fibre Channel over Ethernet (FCoE) network uses FC commands, but transmits them over traditional Ethernet networks. FCoE encapsulates the commands within standard network protocols. In other words, if you have an Ethernet LAN, you can implement an FCoE without all of the overhead costs associated with a traditional Fibre Channel SAN.

Handling Big Data

Big Data refers to data sets that are so large, traditional tools aren't available to analyze them. As an example, Amazon databases are as large as 20 terabytes in size. They access these databases to analyze customer behavior and provide managers with actionable data. Sales managers can ask questions such as what products in any given category do customers look at most often, and what products do these customers buy, and how can we direct more customers to the products that others are buying. The possibilities are endless.

Security personnel use many of the same methods to protect Big Data. This includes using access controls to prevent unauthorized access and encryption to protect the confidentiality of proprietary data.

Understanding Cloud Computing

Cloud computing is one of those catchy terms that has captured our imagination. It just sounds cool. “To the cloud...” You can use cloud-computing technologies to host hardware, such as servers, and to host data.

Even though the phrase *cloud computing* is relatively new, the concept isn’t. In short, cloud computing simply refers to accessing computing resources via a different location than your local computer. In most situations today, you’re accessing these resources through the Internet.

As an example, if you use web-based email such as Gmail, you’re using cloud computing. More specifically, the web-based mail is a Software as a Service cloud computing service. You know that you’re accessing your email via the Internet, but you really don’t know where the physical server hosting your account is located. It could be in a data center in the middle of Virginia, tucked away in Utah, or just about anywhere else in the world.

Cloud computing is very useful for heavily utilized systems and networks. As an example, consider the biggest shopping day in the United States—Black Friday, the day after Thanksgiving, when retailers go into the black. Several years ago, Amazon.com had so much traffic during the Thanksgiving weekend that its servers could barely handle it. The company learned its lesson, though. The next year, it used cloud computing to rent access to servers specifically for the Thanksgiving weekend, and, despite increased sales, it didn’t have any problems.

As many great innovators do, Amazon didn’t look on this situation as a problem, but rather an opportunity. If it needed cloud computing for its heavily utilized system, other companies probably had the same need. Amazon now hosts cloud services to other organizations via its Amazon Elastic Compute Cloud (Amazon EC2) service. Amazon EC2 combines virtualization with cloud computing and they currently provide a wide variety of services via Amazon EC2.

Software as a Service

Software as a Service (SaaS) includes any software or application provided to users over a network such as the Internet. Internet users access the SaaS applications with a web browser. It usually doesn't matter which web browser or operating system a SaaS customer uses. They could be using Internet Explorer, Chrome, Firefox, or just about any web browser.

As mentioned previously, web-based email is an example of SaaS. This includes Gmail, Yahoo! Mail, and others. The service provides all the components of email to users via a simple web browser.

If you have a Gmail account, you can also use Google Docs, another example of SaaS. Google Docs provides access to several SaaS applications, allowing users to open text documents, spreadsheets, presentations, drawings, and PDF files through a web browser.

A talented developer named Lee Graham and I teamed up to create CertApps.com to create study materials. He's an Apple guy running a Mac while I'm a Microsoft guy running Windows, and we live in different states. However, we post and share documents through Google Docs and despite different locations and different applications running on our individual systems, we're able to easily collaborate. One risk is that our data is hosted on Google Docs, and if attackers hack into Google Docs, our data may be compromised.

A specialized version of SaaS is Management as a Service (MaaS). With MaaS, an organization is able to outsource management and monitoring of IT resources. For example, a third party can routinely review logs and provide reports back to the organization.

Multi-tenancy (sometimes referred to as multi-tenant) is a concept associated with cloud computing. A multi-tenancy architecture uses a single instance of an application accessed by multiple customers. You can think of this like a single instance of a web browser accessing multiple web sites in separate tabs. In contrast, single-tenancy architecture creates a separate instance of a SaaS application for each customer. Using the web browser analogy, you'd have a separate web browser window for every site you're visiting. Customer data remains private for customers in both multi-tenancy and single-tenancy architectures.

Platform as a Service

Platform as a Service (PaaS) provides customers with a preconfigured computing platform they can use as needed. It provides the customer with an easy-to-configure operating system, combined with appropriate applications and on-demand computing.

Many cloud providers refer to this as a managed hardware solution. For example, I host <http://gcgapremium.com/> on a virtual server through Liquid Web (<http://www.liquidweb.com/>) using one of their “Fully Managed” offerings.

Liquid Web provides several features in their fully managed solutions, including an installed operating system, a core software package used for web servers, Apache as a web server, antivirus software, spam protection, and more. Additionally, they keep the operating system up to date with relevant updates and patches. I manage the software used for the web site, including software changes and updates. However, I don’t need to worry about managing the server itself. The couple of times when the server developed a problem, they fixed it before I was even aware of the problem.

Infrastructure as a Service

Infrastructure as a Service (IaaS) allows an organization to outsource its equipment requirements, including the hardware and all of its support operations. The IaaS service provider owns the equipment, houses it in its data center, and performs all of the required hardware maintenance. The customer essentially rents access to the equipment and often pays on a per-use basis.

Many cloud providers refer to this as a self-managed solution. They provide access to a server with a default operating system installation, but customers must configure it and install additional software based on their needs. Additionally, customers are responsible for all operating system updates and patches.

IaaS can also be useful if an organization is finding it difficult to manage and maintain servers in its own data center. By outsourcing its requirements, the company limits its hardware footprint. It can do this instead of, or in addition to, virtualizing some of its servers. With IaaS, it needs fewer servers in its data center and fewer resources, such as power, HVAC, and personnel to manage the servers.

Remember this

Applications such as web-based email provided over the Internet are Software as a Service (SaaS) cloud-based technologies. Platform as a Service (PaaS) provides customers with a fully managed platform, which the vendor keeps up to date with current patches. Infrastructure as a Service (IaaS) provides customers with access to hardware in a self-managed platform. Customers are responsible for keeping an IaaS system up to date.

Public Versus Private Cloud

Public cloud services are available from third-party companies. For example, Dropbox and Google operate file-hosting services. Some services are available free and some services cost money. For example, Google offers 15 GB of free storage, but if you want additional storage, you can purchase it from Google.

A private cloud is set up for specific organizations. For example, the Shelbyville Nuclear Power Plant might decide it wants to store data in the cloud, but does not want to use a third-party vendor. Instead, the plant chooses to host its own servers and make these servers available to internal employees through the Internet.

Not all cloud implementations fit exactly into these definitions through. A hybrid cloud is a combination of two or more clouds. They can be all private, all public, or a combination. These retain separate identities to help protect resources in the private cloud. However, they are bridged together, often in such a way that it is transparent to the users.

Cloud Computing Risks

Many cloud-computing solutions are blended. They combine data and operating systems, located in the cloud. Although this can provide benefits because you reduce your own data center's footprint, it also presents risks.

One of the primary drawbacks to cloud computing is that you lose physical control of your data. You often won't even know where the data is stored. Employees at the cloud data center can easily steal your data, and you may not know it until the thief has exploited the data. It's also possible for employees to make mistakes that suddenly grant access to your data to anyone.

As an example, consider a glitch that occurred at Dropbox in June 2011. Dropbox is a cloud storage site used by about 25 million customers to store files such as documents, videos, and photos. After an update to their site, visitors could use any password to access accounts. Obviously, that wasn't the intent of the update. The company's official statement was that the glitch impacted only 1 percent of its customers, but 1 percent of 25 million people is still a lot of people.

Many security professionals say that the only data you should put in the cloud is data you're willing to give away. After all, that's exactly what you may be doing.

Chapter 5 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Implementing Host Security

- Hardening servers and applications makes them more secure from their default installation and improves their overall security posture. Key hardening steps are disabling unnecessary services, protocols, and accounts.
- Disabling unnecessary services and protocols reduces the attack surface. Disabling unnecessary accounts and changing default passwords helps prevent unauthorized access.
- Baselines document the starting point for systems. Security baselines provide a secure starting point for systems and applications. System and application configuration baselines document proper configuration settings. Host software baselines (or application baselines) identify approved software along with software installed on systems.
- Baseline reporting monitors current configurations with a baseline and reports discrepancies. Baselines should be updated when systems are modified to ensure accurate baseline reporting.
- Administrators use baselines to identify anomalies by comparing settings, services, and applications in the baseline with settings, services, and applications on live computers.
- Application whitelisting allows authorized applications to run, but blocks all other applications. Application blacklisting blocks unauthorized applications, but allows other applications to run.
- Virtualization allows multiple servers to operate on a single physical host. They provide increased availability with various tools such as snapshots and easy restoration.
- Sandboxing within virtual environments combined with snapshots provides a high level of flexibility for testing security controls and testing patches.
- Virtual local area networks (VLANs) separate or segment traffic. You create them with physical switches on physical networks and with virtual switches on virtual networks.
- Patch management procedures ensure operating systems and applications are kept up to date with current patches. This ensures they are protected against known vulnerabilities.
- Static computing environments include supervisory control and data acquisition (SCADA) systems and embedded systems. A key method of protecting them is to use a combination of redundant controls and diverse controls. VLANs and network intrusion prevention systems (NIPS) protect networks holding static systems.

Securing Mobile Devices

- You can protect mobile devices with encryption of data, screen locks, and remote wipe capabilities. Remote wiping removes all the data from a lost phone.
- Disabling the use of removable media on mobile devices prevents users from saving data on USB thumb drives and other removable media.
- GPS tracking can help locate lost or stolen mobile devices. Geo-tagging uses GPS features and adds geographical information to files such as pictures when posting them on social media sites.
- Radio-frequency identification (RFID) provides automated inventory control and can detect movement of mobile devices.
- Ensuring mobile devices are up to date with current patches and antivirus signature files is a primary concern. Mobile device management (MDM) tools can ensure that devices meet these requirements and block network access if devices do not meet these requirements.
- Data security is one of the primary security concerns related to bring your own device (BYOD) policies.
- VLANs can isolate mobile device traffic from the primary network.

Protecting Data

- The primary method of protecting the confidentiality of data is with encryption and strong access controls. Encryption protects both data at rest (stored on a device) and data in motion (transmitted over a network).
- You can encrypt individual columns in a database (such as credit card numbers), entire databases, individual files, entire disks, and removable media.
- Whole disk encryption and full device encryption procedures protect all data on a disk and are useful when protecting data on USB flash drives. Many mobile devices and removable devices support full disk encryption. Encryption protects the data if the device is stolen.
- Hardware encryption is faster and more efficient than software encryption.
- A Trusted Platform Module (TPM) is a chip in a motherboard included with many laptops. TPMs have a storage root key used to generate and protect other encryption keys. TPMs support full disk encryption.

- A hardware security module (HSM) is a removable or external device used for encryption. An HSM generates and stores RSA encryption keys and can be integrated with servers to provide hardware encryption.
- Network-based data loss prevention (DLP) devices reduce the risk of data leakage. They can analyze outgoing data, such as emails, and detect when employees send out confidential company data. Endpoint DLP can block users from copying or printing certain files.

Understanding Cloud Computing

- Provider clouds provide increased capabilities for heavily utilized systems and networks.
- Software as a Service (SaaS) includes web-based applications such as web-based email.
- Infrastructure as a Service (IaaS) provides hardware resources via the cloud. It can help an organization limit the size of their hardware footprint and reduce personnel costs.
- Platform as a Service (PaaS) provides an easy-to-configure operating system and on-demand computing for customers.
- Physical control of data is a key security control an organization loses with cloud computing.

Chapter 5 Practice Questions

1. Your organization wants to improve the security posture of internal database servers and protect against zero-day vulnerabilities. Of the following choices, what provides the BEST solution?
 - A. Opening ports on a server's firewall
 - B. Disabling unnecessary services
 - C. Keeping systems up to date with current patches
 - D. Keeping systems up to date with current service packs
2. You need to monitor the security posture of several servers in your organization and keep a security administrator aware of their status. Which of the following tasks will BEST help you meet this goal?
 - A. Establishing baseline reporting
 - B. Determining attack surface
 - C. Implementing patch management
 - D. Enabling sandboxing
3. Maggie is compiling a list of approved software for desktop operating systems within a company. What is the MOST likely purpose of this list?
 - A. Host software baseline
 - B. Baseline reporting
 - C. Application configuration baseline
 - D. Code review
4. Your organization wants to ensure that employees do not install or play operating system games, such as solitaire and FreeCell, on their computers. Which of the following is the BEST choice to prevent this?
 - A. Security policy
 - B. Application whitelisting
 - C. Anti-malware software
 - D. Antivirus software
5. An IT department recently had its hardware budget reduced, but the organization still expects them to maintain availability of services. Of the following choices, what would BEST help them maintain availability with a reduced budget?
 - A. Failover clusters
 - B. Virtualization
 - C. Bollards
 - D. Hashing
6. You are preparing to deploy a new application on a virtual server. The virtual server hosts another server application that employees routinely access. Which of the following is the BEST method to use when deploying the new application?

- A. Take a snapshot of the VM before deploying the new application.
- B. Take a snapshot of the VM after deploying the new application.
- C. Apply blacklisting techniques on the server for the new applications.
- D. Back up the server after installing the new application.

7. A recent risk assessment identified several problems with servers in your organization. They occasionally reboot on their own and the operating systems do not have current security fixes. Administrators have had to rebuild some servers from scratch due to mysterious problems. Which of the following solutions will mitigate these problems?

- A. Virtualization
- B. Sandboxing
- C. IDS
- D. Patch management

8. Administrators ensure server operating systems are updated at least once a month with relevant patches, but they do not track other software updates. Of the following choices, what is the BEST choice to mitigate risks on these servers?

- A. Application change management
- B. Application patch management
- C. Whole disk encryption
- D. Application hardening

9. Homer noticed that several generators within the nuclear power plant have been turning on without user interaction. Security investigators discovered that an unauthorized file was installed and causing these generators to start at timed intervals. Further, they determined this file was installed during a visit by external engineers. What should Homer recommend to mitigate this threat in the future?

- A. Create an internal CA.
- B. Implement WPA2 Enterprise.
- C. Implement patch management processes.
- D. Configure the SCADA within a VLAN.

10. Your company has recently provided mobile devices to several employees. A security manager has expressed concerns related to data saved on these devices. Which of the following would BEST address these concerns?

- A. Disabling the use of removable media
- B. Installing an application that tracks the location of the device
- C. Implementing a BYOD policy
- D. Enabling geo-tagging

11. Which of the following is the MOST likely negative result if administrators do not implement access controls correctly on an encrypted USB hard drive?

- A. Data can be corrupted.
- B. Security controls can be bypassed.
- C. Drives can be geo-tagged.
- D. Data is not encrypted.

12. Your company provides electrical and plumbing services to homeowners. Employees use tablets during service calls to record activity, create invoices, and accept credit card payments. Which of the following would BEST prevent disclosure of customer data if any of these devices are lost or stolen?

- A. Mobile device management
- B. Disabling unused features
- C. Remote wiping
- D. GPS tracking

13. Key personnel in your organization have mobile devices, which store sensitive information. What can you implement to prevent data loss from these devices if a thief steals one?
- A. Asset tracking
 - B. Screen lock
 - C. Mobile device management
 - D. GPS tracking
14. Which of the following represents a primary security concern when authorizing mobile devices on a network?
- A. Cost of the device
 - B. Compatibility
 - C. Virtualization
 - D. Data security
15. Your company is planning on implementing a policy for users so that they can connect their mobile devices to the network. However, management wants to restrict network access for these devices. They should have Internet access and be able to access some internal servers, but management wants to ensure that they do not have access to the primary network where company-owned devices operate. Which of the following will BEST meet this goal?
- A. WPA2 Enterprise
 - B. VPN
 - C. GPS
 - D. VLAN
16. Your organization hosts a web site with a back-end database. The database stores customer data, including credit card numbers. Which of the following is the BEST way to protect the credit card data?
- A. Full database encryption
 - B. Whole disk encryption
 - C. Database column encryption
 - D. File-level encryption
17. Bart copied an encrypted file from his desktop computer to his USB drive and discovered that the copied file isn't encrypted. He asks you what he can do to ensure files he's encrypted remain encrypted when he copies them to a USB drive. What would you recommend as the BEST solution to this problem?
- A. Use file-level encryption.
 - B. Convert the USB to FAT32.
 - C. Use whole disk encryption on the desktop computer.
 - D. Use whole disk encryption on the USB drive.
18. You are comparing different encryption methods. Which method includes a storage root key?
- A. HSM
 - B. NTFS
 - C. VSAN
 - D. TPM
19. Management wants to ensure that employees do not print any documents that include customer PII. Which of the following solutions would meet this goal?
- A. HSM
 - B. TPM
 - C. VLAN
 - D. DLP
20. Of the following choices, which one is a cloud computing option that allows customers to apply patches to the operating system?

- A. Hybrid cloud
- B. Software as a Service
- C. Infrastructure as a Service
- D. Private

Chapter 5 Practice Question Answers

1. **B.** Disabling unnecessary services helps reduce threats, including threats from zero-day vulnerabilities. It also reduces the threat from open ports on a firewall if the associated services are disabled, but opening ports won't reduce threats. Keeping systems up to date with patches and service packs protects against known vulnerabilities and is certainly a good practice. However, by definition, there aren't any patches or service packs available for zero-day vulnerabilities.
2. **A.** Establishing baseline reporting processes allows you to monitor the systems and identify any changes from the baseline that might affect their security posture. You would determine the attack surface prior to establishing a baseline. Patch management is important, but it doesn't monitor the overall security posture of systems. Sandboxing allows you to isolate systems for testing, but isn't used for online production systems.
3. **A.** A host software baseline (also called an application baseline) identifies a list of approved software for systems and compares it with installed applications. Baseline reporting is a process that monitors systems for changes and reports discrepancies. An application configuration baseline identifies proper settings for applications. A code review looks at the actual code of the software, and doesn't just create a list.
4. **B.** Application whitelisting identifies authorized applications and prevents users from installing or running any other applications. Alternately, you can use a blacklist to identify specific applications that cannot be installed or run on a system. A security policy (such as an acceptable use policy) can state a rule to discourage this behavior, but it doesn't enforce the rule by preventing users from installing or running the software. Anti-malware software and antivirus software can detect and block malware, but not applications.
5. **B.** Virtualization provides increased availability because it is much easier to rebuild a virtual server than a physical server after a failure. Virtualization supports a reduced budget because virtual servers require less hardware, less space in a data center, less power, and less heating and air conditioning. Failover clusters are more expensive. Bollards are physical barriers that block vehicles. Hashing provides integrity, not availability.
6. **A.** Taking a snapshot of the virtual machine (VM) before deploying it ensures that the VM can be reverted to the original configuration if the new application causes problems. Taking a snapshot after the installation doesn't allow you to revert the image. Blacklisting prevents an application from running, so it isn't appropriate for a new application deployed on a server. Backing up the server might be appropriate before installing the new application but not after.
7. **D.** Patch management procedures ensure that systems are kept up to date with current security fixes and patches and help eliminate problems with known attack methods. The scenario indicates that these systems have been attacked, exploiting the vulnerabilities caused by not patching them. Virtualization will have the same problems if the systems are not kept up to date. Sandboxing isolates systems for testing, but there isn't any indication these servers should be isolated. An intrusion detection system (IDS) might identify some attacks, but the systems will still be exploited if they aren't patched.
8. **B.** Application patch management practices ensure that applications are kept up to date with relevant patches, similar to how the operating systems are kept up to date with patches. Application change management helps control changes to the applications. Whole disk encryption helps protect confidentiality, but is unrelated to this question. Application hardening secures the applications when they are deployed, but doesn't keep them up to date with current patches.
9. **D.** The generators are likely controlled within a supervisory control and data acquisition (SCADA) system and isolating them within a virtual local area network (VLAN) will protect them from unauthorized access. An internal certificate authority (CA) issues and manages certificates within a Public Key Infrastructure (PKI), but there isn't any indication certificates are in use. Wi-Fi Protected Access II (WPA2) secures wireless networks, but doesn't protect SCADA networks. Patch management processes help ensure systems are kept up to date with patches, but this doesn't apply in this scenario.
10. **A.** Disabling the use of mobile media on the devices will reduce the potential of data loss from these devices. It would make it more difficult to copy data to and from the devices. Tracking the location won't affect data. The devices are provided by the company, so a bring your own device (BYOD) policy isn't relevant. Geo-tagging only refers to geographic location information attached to pictures posted on social media sites.
11. **B.** If access controls are not implemented correctly, an attacker might be able to bypass them and access the data. The incorrect implementation of the access controls won't corrupt the data. Files such as pictures posted on social media can be geo-tagged, but this is unrelated to a hard drive. The scenario says the drive is encrypted, so the data is encrypted.
12. **C.** Remote wiping sends a signal to a device and erases all data, which would prevent disclosure of customer data. Mobile device management helps ensure devices are kept up to date with current patches. Disabling unused features is a basic hardening step for mobile devices, but doesn't help if the device is lost. Global positioning system (GPS) tracking helps locate the device, but doesn't necessarily prevent data disclosure if the device cannot

be retrieved.

13. **B.** A screen lock helps prevent data loss in the event of theft of a mobile device storing sensitive information. Other security controls (not listed as answers in this question) that help prevent loss of data in this situation are account lockouts, full device encryption, and remote wipe capabilities. Asset tracking is an inventory control method. Mobile device management helps keep systems up to date with current patches. Global positioning system (GPS) tracking helps locate the device.

14. **D.** Protecting data is a primary security concern when authorizing mobile devices on a network, often because mobile devices are more difficult to manage. The cost of the devices is trivial when compared with the cost of other network devices and the value of data. Compatibility issues aren't a major concern and typically only affect the ability to use an application. Virtualization techniques can be used with mobile devices allowing users to access virtual desktops, but these enhance security.

15. **D.** A virtual local area network (VLAN) provides network segmentation and can prevent employee owned devices from accessing the primary network. WPA2 Enterprise provides strong security for the devices by ensuring they authenticate through an 802.1x server, but this doesn't segment them on a separate network. A virtual private network (VPN) allows remote employees to connect to a private network, but is unrelated to this question. A global positioning system (GPS) is useful for locating lost devices but not segmenting network traffic.

16. **C.** Database column (or field) encryption is the best choice because it can be used to encrypt the fields holding credit card data, but not fields that don't need to be encrypted. Full database encryption and whole disk encryption aren't appropriate because everything doesn't need to be encrypted to protect the credit card data. File-level encryption isn't appropriate on a database and will often make it inaccessible to the database application.

17. **D.** The best solution is to use whole disk encryption on the USB drive. The scenario indicates Bart is using file-level encryption (such as NTFS encryption) on the desktop computer, but the USB drive doesn't support it, possibly because it's formatted as a FAT32 drive. The result is that the system decrypts the file before copying it to the USB drive. Another solution is to convert the USB to NTFS. Whole disk encryption on the desktop computer wouldn't protect files copied to the USB drive.

18. **D.** A Trusted Platform Module (TPM) includes a storage root key. The TPM generates this key when a user activates the TPM. A hardware security module (HSM) uses RSA keys, but not a storage root key. NT File System (NTFS) supports encryption with Encrypting File System (EFS). A virtual storage area network (VSAN) is a virtualization technique, and it doesn't provide encryption.

19. **D.** A data loss prevention (DLP) solution can limit documents sent to a printer to be printed using content filters. A hardware security module (HSM) and a Trusted Platform Module (TPM) both provide full disk encryption, but cannot block documents sent to a printer. A virtual local area network (VLAN) segments traffic, but isn't selective about documents sent to a printer.

20. **C.** Infrastructure as a Service (IaaS) is a cloud computing option where the vendor provides access to a computer, but customers must manage the system, including keeping it up to date with current patches. A hybrid cloud is a combination of a public cloud and a private cloud. Software as a Service (SaaS) provides access to applications, such as email. An IaaS solution can be public, private, or a hybrid solution.

Chapter 6

Understanding Malware and Social Engineering

CompTIA Security+ objectives covered in this chapter:

1.1 Implement security configuration parameters on network devices and technologies.

- Spam filter

2.6 Explain the importance of security related awareness and training.

- User habits (Prevent tailgating)
- New threats and new security trends/alerts (New viruses, Phishing attacks, Zero-day exploits)

3.1 Explain types of malware.

- Adware, Virus, Spyware, Trojan, Rootkits, Backdoors, Logic bomb, Botnets, Ransomware, Polymorphic malware, Armored virus

3.2 Summarize various types of attacks.

- Spam, Phishing, Spim, Vishing, Spear phishing, Privilege escalation

3.3 Summarize social engineering attacks and the associated effectiveness with each attack.

- Shoulder surfing, Dumpster diving, Tailgating, Impersonation, Hoaxes, Whaling, Vishing
- Principles (reasons for effectiveness) (Authority, Intimidation, Consensus/Social proof, Scarcity, Urgency, Familiarity/liking, Trust)

4.3 Given a scenario, select the appropriate solution to establish host security.

- Anti-malware (Antivirus, Anti-spam, Anti-spyware, Pop-up blockers)

Malicious software (malware) and social engineering are two common threats that any organization will face. Within information technology (IT) security, these are relatively easy to prevent. However, that doesn't reduce their importance. The damage can be extensive if an organization ignores these threats.

Understanding Malware Types

Malware includes a wide range of software that has malicious intent. Malware is not software that you would knowingly purchase or download and install. Instead, it is installed onto your system through devious means. Infected systems give various symptoms, such as running slower, starting unknown processes, sending out email without user action, random reboots, and more.

You might hear people use the term *virus* to describe all types of malware, but that isn't accurate. A virus is a specific type of malware, and malware includes many other types of malicious software, including worms, logic bombs, Trojans, ransomware, rootkits, spyware, and more.

Viruses

A *virus* is a set of malicious code that attaches itself to a host application. The host application must be executed to run, and the malicious code executes when the host application is executed. The virus tries to replicate by finding other host applications to infect with the malicious code. At some point, the virus activates and delivers its payload.

Typically, the payload of a virus is damaging. It may delete files, cause random reboots, join the computer to a botnet, or enable backdoors that attackers can use to access systems remotely. Some older viruses merely displayed a message at some point, such as “Legalize Marijuana!” Most viruses won’t cause damage immediately. Instead, they give the virus time to replicate first.

A user will often execute the virus (though unknowingly), but other times, an operating system will automatically execute it after user interaction. For example, when a user plugs in an infected USB drive, the system can execute the virus infecting the system. Note that not all malware needs user interaction to run. As an example, worms are self-replicating and do not need user interaction.

Operation Buckshot Yankee

William Lynn, a U.S. Deputy Secretary of Defense, wrote an article in the *Foreign Affairs* magazine that demonstrates the risk from USB drives. He indicated that this incident marked a turning point in the U.S. cyber defense strategy.

In 2008, the U.S. military suffered a significant data breach that they traced back to a USB flash drive. Apparently, a foreign intelligence agency developed malware and installed it on a USB drive. Someone, though no one seems to be saying who, inserted the USB drive into a military laptop somewhere in the Middle East. The malware quickly infected the system.

The malware continued to operate silently on the mobile system and ultimately infected the U.S. Central Command's network, including both classified and unclassified systems. Reports indicate that attackers were able to transfer data from the network to foreign servers. Ultimately, the U.S. military discovered the malware and launched Operation Buckshot Yankee. They cleaned the virus off all systems and investigated the incident. It is clear that this was a major incident, even though it started from malware on a single USB drive.

I was working on a U.S. base in 2008 when a new rule came out that banned the use of all removable USB flash drives. There was no mention of Operation Buckshot Yankee at the time, but it was clear that they were serious about enforcing the rule. I know of one contractor who ignored the rule, plugged in a USB flash drive, and had an opportunity to upgrade his résumé the next day. He was fired.

. . .

Armored Virus

When antivirus (AV) researchers discover a new virus, they typically attempt to reverse engineer the code. Application developers first write applications in a computer language, such as C, C++, and C#. Although these have specific syntax rules, they are easy to read by people who know the language. Developers then compile the code into an executable application. Reverse engineering code is the process of decompiling the executable application and analyzing the code to discover what it does.

Armored viruses use various techniques to make the reverse engineering process more difficult for the AV researchers. Some methods used by armored viruses are:

- **Complex code.** Some armored viruses use confusing code specifically designed to mask what the virus is actually trying to do.
- **Encryption.** Some compilers encrypt the code with the virus, making it more difficult to decompile. This code must first be decrypted before it can be decompiled.
- **Hiding.** Some viruses attempt to hide their actual location by tricking AV software into thinking the file is located somewhere else.

Remember this

An *armored virus* uses one or more techniques to make it difficult to reverse engineer. Common techniques include using complex code, using encryption, or hiding the location.

Polymorphic Malware

Some virus developers use polymorphism as a method of armor. *Polymorphic* malware has the ability to morph or mutate when it replicates itself, or when it executes. The goal is to create a virus or other malware with enough variations that AV software cannot detect it as the same malware. Over time, a single malware file could have thousands of variants.

Virus developers typically encrypt polymorphic viruses as an additional layer of armor to evade detection by AV software. The virus includes a decryption method that executes when the file executes. A common polymorphic technique is to vary the encryption/decryption method slightly. This method retains the original malicious code, but modifies the file enough to make it difficult to detect.

Worms

A *worm* is self-replicating malware that travels throughout a network without the assistance of a host application or user interaction. A worm resides in memory and is able to use different transport protocols to travel over the network.

One of the significant problems caused by worms is that they consume network bandwidth. Worms can replicate themselves hundreds of times and spread to all the systems in the network. Each infected system tries to locate and infect other systems on the network, and network performance can slow to a crawl.

Logic Bombs

A *logic bomb* is a string of code embedded into an application or script that will execute in response to an event. The event may be a specific date or time, when a user launches a specific program, or any event the programmer decides on.

There's an often-repeated story about a company that decided it had to lay off an engineer due to an economic downturn. His bosses didn't see him doing much, so they thought they could do without him. Within a couple of weeks after he left, they started having all sorts of computer problems they just couldn't resolve.

They called him back, and within a couple of weeks, everything was fine. A few months later, they determined they had to lay him off again. You guessed it. Within a couple of weeks, things went haywire again.

The engineer had programmed a logic bomb that executed when the payroll program ran. It checked for his name on the payroll, and when it was there, things were fine, but when his name wasn't there, ka-boom!—the logic bomb exploded.

Remember this

A logic bomb executes in response to an event, such as when a specific application is executed or a specific time arrives.

Backdoors

A *backdoor* provides another way of accessing a system, similar to how a backdoor in a house provides another method of entry. Malware such as Trojans often install backdoors on systems to bypass normal authentication methods.

Application developers often code backdoors into applications, but this practice is not recommended. For example, an application developer might create a backdoor within an application intended for maintenance purposes. However, if attackers discover the backdoor, they can use it to access the application.

Effective account management policies help prevent ex-employees from creating backdoors after they are fired. For example, if an employee loses network access immediately after being fired, the employee cannot create a backdoor account. In contrast, if an administrator retains network access, he has the opportunity to create another administrative account that he can use even if his original account is disabled. That's exactly what a Fannie Mae Unix engineer did after being told he was fired.

Fannie Mae's account management policy did not revoke his elevated system privileges right away, giving him time to create a backdoor account. After going home, he accessed the system remotely and installed a logic bomb script scheduled to run at 9:00 a.m. on January 31. If another administrator hadn't discovered the logic bomb, it would have deleted data and backups for about four thousand servers, changed their passwords, and shut them down.

Remember this

A backdoor provides another of way of accessing a system. Many types of malware create backdoors, allowing attackers to access systems from remote locations. Employees have also created backdoors in applications and systems.

Trojans

A *Trojan*, also called a Trojan horse, looks like something beneficial, but it's actually something malicious. According to a report by PandaLabs, Trojans represented over 70 percent of new malware strains in 2013, and they also represented 78 percent of malware infections. Trojans frequently create backdoors, allowing criminals to connect to systems remotely after they are infected.

Trojan horses are named after the infamous horse from the Trojan War. In Greek mythology, the Achaeans tried to sack the city of Troy for several years, but they simply couldn't penetrate the city's defenses. At some point, someone got the idea of building a huge wooden horse and convincing the people of Troy that it was a gift from the gods. Warriors hid inside, and the horse was rolled up to the gates.

The people of Troy partied all day and all night celebrating their good fortune, but when the city slept, the warriors climbed down from inside the horse and opened the gates. The rest of the warriors flooded in. What the Greek warriors couldn't do for years, the Trojan horse helped them do in a single day.

In computers, a Trojan horse can come as pirated software, a cool screen saver, a useful utility, a game, or something else that users may be enticed to download and try. Attackers are increasingly using drive-by downloads to deliver Trojans. In a *drive-by download*, web servers include malicious code that attempts to download and install itself on user computers after the user visits. Here are the typical steps involved in a drive-by download:

1. Attackers compromise a web site to gain control of it.
2. Attackers install a Trojan embedded in the web site's code.
3. Attackers attempt to trick users into visiting the site. Sometimes, they simply send the link to thousands of users via email hoping that some of them click the link.
4. When users visit, the web site attempts to download the Trojan onto the users' systems.

Another Trojan method that has become popular in recent years is rogueware, also known as scareware. *Rogueware* masquerades as a free antivirus program. When a user visits a site, a message on the web page or a pop-up appears indicating it detected malware on the user's system. The user is encouraged to download and install free antivirus software.

On the surface, this free antivirus software looks useful. However, it isn't. If a user installs and runs it on a system, it appears to do a system scan. After the scan completes, it reports finding multiple issues, such as infections by dozens of viruses. The report isn't true. The application reports these issues even on a freshly installed operating system with zero infections.

It then encourages the user to resolve these issues immediately. If the user tries to resolve the issues, the program informs the user that this is only the trial version, and the trial version won't resolve these issues. However, for the small fee of \$79.95, users can unlock the full version to remove the threats. Some rogueware installs additional malicious components. For example, it can join the computer to a botnet, or allow the attacker to take remote control of the infected system.

As mentioned previously, you can also infect a system by plugging in an infected USB flash drive. In this case, the attacker can install the Trojan onto several USB drives and leave them lying around. Someone picks one up, plugs it in, and the system is infected. The system then infects other USBs, which infect other systems.

Remember this

A Trojan appears to be something useful but includes a malicious component, such as installing a backdoor on a user's system. Many Trojans are delivered via drive-by downloads. They can also infect systems from rogueware, pirated software, games, or infected USB drives.

Botnets

A *botnet* combines the words *robot* and *network*. It includes multiple computers that act as software robots and function together in a network (such as the Internet), often for malicious purposes. The computers in a botnet are called *zombies* and they will do the bidding of whoever controls the botnet.

Bot herders are criminals who manage botnets. They attempt to infect as many computers as possible and control them through one or more servers running command-and-control software. The infected computers periodically check in with the command-and-control servers, receive direction, and then go to work. The user is often unaware of the activity.

Most computers join a botnet through malware infection. For example, a user could download pirated software with a Trojan or click a malicious link, resulting in a drive-by download. The malware then joins the system to a botnet.

As an example, Coreflood malware is a Trojan horse that opens a backdoor on compromised computers. Authorities shut down the Coreflood botnet in April 2011, and its command-and-control servers managed about 2.3 million computers at that time. Experts estimate they had stolen between \$10 million and \$100 million before authorities shut them down.

Infecting 2.3 million computers and stealing tens of millions of dollars draws a lot of attention. To avoid attention, many botnets manage fewer than 50,000 computers and fly under the radar of most authorities. The result is the same for the victims, though. It doesn't matter if victims are robbed by a huge botnet or a smaller botnet; they have still been robbed.

Botnet herders sometimes maintain complete control over their botnets. Other times, they rent access out to others to use as desired. Some of the instructions sent by the command-and-control servers include:

- Send spam.
- Launch a distributed denial-of-service attack.
- Download additional malware, adware, or spyware such as keyloggers.

Ransomware

A specific type of Trojan is ransomware. Attackers take control of the user's computer and then demand the user pay a ransom to get the control back. Criminals often deliver ransomware via drive-by downloads or embedded in other software. Two ransomware viruses that have attacked many people are the Police Virus and CryptoLocker, and they provide good examples of how ransomware works.

The Police Virus

The Police Virus (also known as Trojan Reveton, Police Ukash, and Moneyapak virus) accuses users of being involved in illegal activities and demands they pay a fine. It often displays a notification from a law enforcement agency such as the U.S. FBI, the Australian Federal Police, or the Metropolitan Police when the computer boots. In some cases, it takes control of the webcam and displays activity in the user's room.

It typically demands victims pay \$100 or €100, depending on their location. In some cases, the Police Virus demands as much as \$300 or €300 as a fine or penalty. It promises to remove the messages and return full control of the computer to the user, if the user pays. One piece of good news is that the Police Virus doesn't actually encrypt or destroy any data. That isn't the case with CryptoLocker.

CryptoLocker

CryptoLocker doesn't try to trick the user, but instead uses basic kidnapping and ransom tactics. For example, many kidnappers have abducted a child, and then attempted to get a ransom from a parent to release the child. CryptoLocker doesn't abduct people, but it does take control of valuable user files.

After CryptoLocker takes control of the user's computer, it encrypts valuable user files, such as photos, videos, and text documents. It then demands the user pay a ransom of up to \$300 or €300. It typically displays a message indicating that the criminals will destroy the decryption key in 72 hours if the user doesn't pay, effectively locking the user's data forever. In some cases, it shows a timer counting down to zero, adding a sense of urgency to the victim to pay quickly.

Because CryptoLocker uses strong asymmetric encryption techniques to encrypt valuable user files, it is almost impossible to decrypt the data in any reasonable amount of time. In addition to encrypting data on the user's computer, it also searches for any network drives and encrypts files on them, too.

As another twist, the CryptoLocker criminals have recently been offering to restore data after the 72 hours expired, at a highly inflated price of as much as \$2,300. It indicates the criminals have saved the decryption key, but removed its association with the victim. If the victim gives them an encrypted file, they'll use it to discover the original encryption key.

Just as many parents are willing to pay ransom to save their children, the success of these types of ransomware indicate many people are willing to pay ransom to restore their data. PandaLabs reported in their 2013 annual report that ransomware has been on the rise and is one of the most common types of malware. They predicted it will be one of the most pervasive threats in 2014.

Remember this

Ransomware is a type of malware that takes control of a user's system or data. Criminals then attempt to extort payment from the victim. Ransomware often includes threats of damaging a user's system or data if the victim does not pay the ransom.

Rootkits

A *rootkit* is a group of programs (or, in rare instances, a single program) that hides the fact that the system has been infected or compromised by malicious code. A user may suspect something is wrong, but antivirus scans and other checks may indicate everything is fine because the rootkit hides its running processes to avoid detection.

In addition to modifying the internal operating system processes, rootkits often modify system files such as the Registry. In some cases, the rootkit modifies system access, such as removing users' administrative access.

Rootkits have system-level access to systems. This is sometimes called root-level access, or kernel-level access, indicating that they have the same level of access as the operating system. Rootkits use hooked processes, or hooking techniques, to intercept calls to the operating system. In this context, *hooking* refers to intercepting system-level function calls, events, or messages. The rootkit installs the hooks into memory and uses them to control the system's behavior.

Antivirus software often makes calls to the operating system that could detect malware, but the rootkit prevents the antivirus software from making these calls. This is why antivirus software will sometimes report everything is OK, even if the system is infected with a rootkit. However, antivirus software can often detect the hooked processes by examining the contents of the system's random access memory (RAM).

Another method used to detect rootkits is to boot into safe mode, or have the system scanned before it boots, but this isn't always successful. It's important to remember that rootkits are very difficult to detect because they can hide so much of their activity. A clean bill of health by a malware scanner may not be valid.

The Trojan.Popureb/E rootkit is an example of a rootkit. Among other things, it overwrites the hard drive's Master Boot Record (MBR), where code is stored to start the operating system. The code on the MBR starts before the operating system boots and it remains invisible to the operating system and security software. Even when antivirus software detects the rootkit, the rootkit protects itself. It prevents any attempts to overwrite the MBR by changing write operations to read operations, though it reports that the write operation completed successfully.

It's important to remember that behind any type of malware, you'll likely find an attacker involved in criminal activity. Attackers who have successfully installed a rootkit on a user's system might log on to the user's computer remotely, using a backdoor installed by the rootkit. Similarly, attackers might direct the computer to connect to computers on the Internet and send data. Data can include anything collected from a keylogger, collected passwords, or specific files or file types stored on the user's computer.

Remember this

Rootkits have system-level or kernel-level access and can modify system files and system access. Rootkits hide their running processes to avoid detection with hooking techniques. Tools that can inspect RAM can discover these hidden hooked processes .

Spyware

Spyware is software installed on users' systems without their awareness or consent. Its purpose is often to monitor the user's computer and the user's activity. Spyware takes some level of control over the user's computer to learn information and sends this information to a third party. If spyware can access a user's private data, it results in a loss of confidentiality.

Some examples of spyware activity are changing a user's home page, redirecting web browsers, and installing additional software, such as search engines. In some situations, these changes can slow a system down, resulting in poorer performance. These examples are rather harmless compared with what more malicious spyware (called privacy-invasive software) may do.

Privacy-invasive software tries to separate users from their money using data-harvesting techniques. It attempts to gather information to impersonate users, empty bank accounts, and steal identities. For example, some spyware includes keyloggers used to capture keystrokes. The keystrokes are stored in a file, and the spyware periodically sends the file to the attacker. In some instances, the spyware allows the attacker to take control of the user's system remotely.

Spyware is often included with other software like a Trojan. The user installs one application but unknowingly gets some extras. Spyware can also infect a system in a drive-by download. The user simply visits a malicious web site that includes code to automatically download and install the spyware onto the user's system.

Adware

When *adware* first emerged, its intent was usually to learn a user's habits for the purpose of targeted advertising. As the practice of gathering information on users became more malicious, more people began to call it spyware. However, some traditional adware still exists.

A common type of adware is pop-ups. For example, while you are visiting a site, another browser window appears, or pops up, with an advertisement. These pop-up windows aren't malicious, but they are annoying.

Sometimes pop-ups can be helpful. As a legitimate example, my online bank has interest-rate information that I can view. When I click on this link, it pops up another window showing the interest-rate information without taking me away from the current page I'm viewing.

The term adware also applies to software that is free but includes advertisements. The user is well aware that the advertisements appear, and has the option to purchase a version of the software that does not include the ads. All of this is aboveboard without any intention of misleading the user.

Remember this

Spyware monitors a user's computer. Pop-ups are annoying windows that appear while browsing. Many pop-ups are adware designed to market products to users.

Recognizing Common Attacks

In addition to malware, it's important to understand some other common attacks. Social engineering includes several techniques attackers use to trick users. Additionally, many attackers use email, instant messaging, and the phone to deliver attacks.

Social Engineering

Social engineering is the practice of using social tactics to gain information. It's often low-tech and encourages individuals to do something they wouldn't normally do, or causes them to reveal some piece of information, such as user credentials. Some of the individual methods and techniques include:

- Flattery and conning
- Assuming a position of authority
- Encouraging someone to perform a risky action
- Encouraging someone to reveal sensitive information
- Impersonating someone, such as an authorized technician
- Tailgating or closely following authorized personnel without providing credentials

In the movie *Catch Me If You Can*, Leonardo DiCaprio played Frank Abagnale Jr., an effective con artist. He learned some deep secrets about different professions by conning and flattering people into telling him. He then combined all he learned to impersonate pilots and doctors and perform some sophisticated forgery.

Social engineers con people in person, as Frank Abagnale Jr. did, and they use other methods as well. They may use the phone, send email with phishing tactics, and even use some trickery on web sites, such as fooling someone into installing rogueware.

As an example of a social engineer using the phone, consider this scenario. Maggie is busy working and receives a call from hacker Herman, who identifies himself as a member of the IT department.

Hacker Herman: "Hi, Maggie. I just wanted to remind you, we'll be taking your computer down for the upgrade today, and it'll be down for a few hours."

Maggie: "Wait. I didn't hear anything about this. I need my computer to finish a project today."

Hacker Herman: "You should have gotten the email. I'm sorry, but I have to get the last few computers updated today."

Maggie: "Isn't there any other way? I really need my computer."

Hacker Herman: "Well...it is possible to upgrade it over the network while you're still working. We don't normally do it that way because we need the user's password to do it."

Maggie: "If I can still work on my computer, please do it that way."

Hacker Herman: "OK, Maggie. Don't tell anyone I'm doing this for you, but if you give me your username and password, I'll do this over the network."

This is certainly a realistic scenario, and many end users will give out their passwords unless security-related awareness and training programs repeatedly repeat the mantra: "Never give out your password."

Attackers aren't always so blatant though. Many times, instead of asking you for your password outright, they ask questions they can use in a password reset system to reset your password. A skilled con man can ask these questions as though he's generally interested in you. Before you know it, you've told him the name of your first dog, your childhood best friend, the name of your first boss, and more. When people post this information in social media, attackers don't even need to ask.

Remember this

Social engineering uses social tactics to trick users into giving up information or performing actions they wouldn't normally take. Social engineering attacks can occur in person, over the phone, while surfing the Internet, and via email.

Impersonating

Some social engineers often attempt to impersonate others. The goal is to convince an authorized user to provide some information, or help the attacker defeat a security control.

As an example, an attacker can impersonate a repair technician to gain access to a server room or telecommunications closet. After gaining access, the attacker can install hardware such as a rogue access point to capture data and send it wirelessly to an outside collection point. Identity verification methods are useful to prevent the success of impersonation attacks. Similarly, attackers impersonate legitimate organizations over the phone and try to gain information.

Shoulder Surfing

Shoulder surfing is simply looking over the shoulder of someone to gain information. The goal is to gain unauthorized information by casual observation, and it's likely to occur within an office environment. This can be to learn credentials, such as a username and password, or a PIN used for a smart card or debit card. Recently, attackers have been using cameras to monitor locations where users enter PINs, such as at automatic teller machines

(ATMs).

A simple way to prevent shoulder surfing is to position monitors and other types of screens so that unauthorized personnel cannot see them. This includes ensuring people can't view them by looking through a window or from reception areas.

Remember this

A social engineer can gain unauthorized information just by looking over someone's shoulder. This might be in person, such as when a user is at a computer, or remotely using a camera.

Tricking Users with Hoaxes

A *hoax* is a message, often circulated through email, that tells of impending doom from a virus or other security threat that simply doesn't exist. Users may be encouraged to delete files or change their system configuration.

An older example is the teddy bear virus (*jdbgmgr.exe*), which was not a virus at all. Victims received an email saying this virus lies in a sleeping state for 14 days and then it will destroy the user's system. It then told users that they can protect their system by deleting the file (which has an icon of a little bear), and provided instructions on how to do so. Users who deleted the file lost some system capability.

More serious virus hoaxes have the potential to be as damaging as a real virus. If users are convinced to delete important files, they may make their systems unusable. Additionally, they waste help-desk personnel's time due to needless calls about the hoax or support calls if users damaged their systems in response to the hoax.

Tailgating and Mantraps

Tailgating is the practice of one person following closely behind another without showing credentials. For example, if Homer uses a badge to gain access to a secure building and Francesca follows closely behind Homer without using a badge, Francesca is tailgating.

Employees often do this as a matter of convenience and courtesy. Instead of shutting the door on the person following closely behind, they often hold the door open for the person. However, this bypasses the access control, and if employees tailgate, it's very easy for a nonemployee to slip in behind someone else. Often, all it takes is a friendly smile from someone like Francesca to encourage Homer to keep the door open for her.

Chapter 2, "Exploring Control Types and Methods," introduced tailgating in the context of physical security controls. As a reminder, mantraps and security guards provide the best prevention for tailgating. A simple mantrap can be a turnstile similar to those used in subways or bus stations. Imagine two men trying to go through a turnstile like this together. It's just not likely. Security guards can check the credentials of each person, and they won't be fooled by a smile as easily. Cameras with a recording capability provide a cheaper alternative.

Dumpster Diving

Dumpster diving is the practice of searching through trash or recycling containers to gain information from discarded documents. Many organizations either shred or burn paper instead of throwing it away.

For example, old copies of company directories can be valuable to attackers. They may identify the names, phone numbers, and titles of key people within the organization. Attackers may be able to use this information in a whaling attack against executives or social engineering attacks against anyone in the organization. An attacker can exploit any document that contains detailed employee or customer information, and can often find value in seemingly useless printouts and notes.

On a personal basis, preapproved credit applications or blank checks issued by credit card companies can be quite valuable to someone attempting to gain money or steal identities. Documentation with any type of Personally Identifiable Information (PII) should be shredded or burned.

Remember this

Dumpster divers search through trash looking for information. Shredding or burning papers instead of throwing them away mitigates this threat.

Recognizing Other Attacks

Beyond social engineering, users should know about many common attacks. This includes attacks such as phishing, spear phishing, whaling, and vishing.

Spam

Spam is unwanted or unsolicited email. Depending on which study you quote, between 80 percent and 92 percent of all Internet email is spam. Some spam is harmless advertisements, while much more is malicious. Spam can include malicious links, malicious code, or malicious attachments. Even when it's not malicious, when only 1 of 10 emails is valid, it can waste a lot of your time.

In some cases, legitimate companies encourage users to opt-in to their email lists and then send them email about their products. When users opt-in to a mailing list, they agree to the terms. On the surface, you'd think that this means that you agree to receive email from the company and that's true. However, terms often include agreeing to allow their partners to send you email, which means the original company can share your email address with others.

Legitimate companies don't send you malicious spam, but they might send you more email than you want. Laws require them to include the ability to opt-out, indicating you don't want to receive any more emails from them. Once you opt-out, you shouldn't receive any more emails from that company.

Originally, spam was just unwanted advertisements sent out to people, even if they chose to opt-out. However, attackers began using spam for malicious purposes. Attackers often include malicious attachments and malicious code within spam email. More recently, spam attacks include malicious links. If users click on a link in a malicious email, it often takes them to a site hosting a drive-by download, as mentioned in the earlier "Trojan" section.

Criminals use a variety of methods to collect email addresses. They buy lists from other criminals and harvest them from web sites. Some malware scans address books of infected computers to collect email. Because they are criminals, they don't care about laws, but they might include opt-out instructions in spam they send. However, instead of this getting you off the email list, it provides them confirmation that your email address is valid. The result is more spam.

Phishing

Phishing is the practice of sending email to users with the purpose of tricking them into revealing personal information or clicking on a link. A phishing attack often sends the user to a malicious web site that appears to the user as a legitimate site.

The classic example is where a user receives an email that looks like it came from eBay, PayPal, a bank, or some other well-known company. The "phisher" doesn't know if the recipient has an account at the company, just as a fisherman doesn't know if any fish are in the water where he casts his line. However, if the attacker sends out enough emails, the odds are good that someone who receives the email has an account.

The email may look like this:

"We have noticed suspicious activity on your account. To protect your privacy, we will suspend your account unless you are able to log in and validate your credentials. Click here to validate your account and prevent it from being locked out."

The email often includes the same graphics that you would find on the vendor's web site or an actual email from the vendor. Although it might look genuine, it simply isn't. Legitimate companies do not ask you to revalidate your credentials via email. If you go directly to the site, you might be asked to provide additional information to prove your identity beyond your credentials, but legitimate companies don't send emails asking you to follow a link and input your credentials to validate them.

Remember this

Spam is unwanted email. Phishing is malicious spam. Attackers attempt to trick users into revealing sensitive or personal information or clicking on a link. Links within email can also lead unsuspecting users to install malware.

Beware of Email from Friends

Criminals have become adept at impersonating your friends. They scan social media sites and identify your friends and family. They then send emails to you that look like they are from your friends or family members, but they really aren't.

As an example, imagine you're friends with Lisa Simpson. You might receive an email that includes "Lisa Simpson" in the From block. However, if you look at the actual email address, you'd find it is something different, such as xyz@yahoo.com. Many times, the owner of the email account (xyz@yahoo.com in this example) is unaware that criminals are using it to send malicious email. A common scenario is that the email account owner's computer is part of a botnet, and the bot herder is sending spam and other malicious email to people via this email account.

Phishing to Install Malware

One phishing email looked like it was from a news organization with headlines of recent news events. If the user clicked anywhere in the email, it showed a dialog box indicating that the user’s version of Adobe Flash was too old to view the story. It then asked, “Would you like to upgrade your version of Adobe Flash?” If the user clicked Yes, it downloaded and installed malware.

Another email had the subject line “We have hijacked your baby” and the following content:

“You must pay once to us \$50,000. The details we will send later. We have attached photo of your family.”

The English seems off, and the receiver may not even have a baby, making this look bogus right away. However, the attackers are only trying to pique your curiosity. If a user clicks on a photo to look at it, it installs malware on the user’s system.

Phishing to Validate Email Addresses

A simple method used to validate email addresses is the use of beacons. A *beacon* is a link included in the email that links to an image stored on an Internet server. The link includes unique code that identifies the receiver’s email address.

For the email application to display the image, it must retrieve the image from the Internet server. When the server hosting the image receives the request, it logs the user’s email address, indicating it’s valid. This is one of the reasons that most email programs won’t display images by default.

Phishing to Get Money

The classic Nigerian scam (also called a 419 scam) is alive and well. You receive an email from someone claiming a relative or someone else has millions of dollars. Unfortunately, the sender can’t get the money without your help. The email says that if you help retrieve the money, you’ll get a substantial portion of the money for your troubles.

This scam often requires the victim to pay a small sum of money with the promise of a large sum of money. However, the large sum never appears. Instead, the attackers come up with reasons why they need just a little more money. In many cases, the scammers request access to your bank account to deposit your share, but instead they use it to empty your bank account.

There are countless variations. Lottery scams inform email recipients they won. Victims sometimes have to pay small fees to release the funds or provide bank information to get the money deposited. They soon learn there is no prize.

In the past, criminals and criminal organizations were behind these frauds. However, many terrorist organizations have been using these types of frauds to raise money. For example, Ultrascan Advanced Global Investigations reports that these types of frauds are the primary funding vehicle for African-based terror groups such as Boko Haram. Boko Haram kidnapped more than 200 schoolgirls from a dormitory in April 2014, and kidnapped 60 more women and children in June 2014.

Spear Phishing

Spear phishing is a targeted form of phishing. Instead of sending the email out to everyone indiscriminately, a spear phishing attack attempts to target specific groups of users, or even a single user. Spear phishing attacks may target employees within a company or customers of a company.

As an example, an attacker might try to impersonate the CEO of an organization in an email. It’s relatively simple to change the header of an email so that the From field includes anything, including the CEO’s name and title. Attackers can send an email to all employees requesting that they reply with their password. Because the email looks like it’s coming from the CEO, these types of phishing emails fool many users.

One solution that deters the success of these types of spear phishing attacks is to use digital signatures. The CEO and anyone else in the company can sign their emails with a digital signature. This provides a high level of certainty to personnel on who sent the email. Chapter 10, “Understanding Cryptography,” covers digital signatures in great depth.

Whaling

Whaling is a form of spear phishing that attempts to target high-level executives. Las Vegas casinos refer to the big spenders as whales, and casino managers are willing to spend extra time and effort to bring them into their casinos. Similarly, attackers consider high-level executives the whales, and attackers are willing to put in some extra effort to catch a whale because the payoff can be so great. When successful, attackers gain confidential company information that they might not be able to get anywhere else.

As an example, attackers singled out as many as 20,000 senior corporate executives in a fine-tuned phishing attack. The emails looked like official subpoenas requiring the recipient to appear before a federal grand jury and included the executive’s full name and other details, such as their company name and phone number. The emails also included a link for more details about the subpoena. If the executives clicked the link, it took them to a web site that indicated they needed a browser add-on to read the document. If they approved this install, they actually installed a keylogger and malware. The keylogger recorded all their keystrokes to a file, and the malware gave the attackers remote access to the executives’ systems.

Similar whaling attacks have masqueraded as complaints from the Better Business Bureau or the Justice Department. Executives are sensitive to issues that may affect the company’s profit, and these attacks get their attention. Although not as common, some whaling attacks attempt to reach the executive via phone to get the data. However, many executives have assistants who screen calls to prevent attackers from reaching the executive via

phone.

Remember this

A spear phishing attack targets specific groups of users. It could target employees within a company or customers of a company. Digital signatures provide assurances to recipients about who sent an email, and can reduce the success of spear phishing. Whaling targets high-level executives.

Spim

Spim is a form of spam using instant messaging (IM). It targets instant messaging users, such as those using Yahoo! Messenger or Windows Live Messenger. Some social media sites also support pop-up windows, and attackers sometimes try to impersonate a user's friend and encourage the victim to click on a link. Many IM services support the use of whitelists, where users identify who they'll receive messages from and the IM system blocks all other messages.

Vishing

Vishing attacks use the phone system to trick users into giving up personal and financial information. It often uses Voice over IP (VoIP) technology and tries to trick the user similar to other phishing attacks. When the attack uses VoIP, it can spoof caller ID, making it appear as though the call came from a real company.

In one form, a machine leaves a phone message saying that you need to return the call concerning one of your credit cards. In another form, you receive an email with the same information. If you call, you'll hear an automated recording giving some vague excuse about a policy and prompting you to verify your identity. One by one, the recording prompts you for more information, such as your name, birthday, Social Security number, credit card number, expiration date, and so on. Sometimes, the recording asks for usernames and passwords. If you give all the requested information, the recording indicates they have verified your account. In reality, you just gave up valuable information on yourself.

Another example of vishing is a just a regular phone call from a criminal. A popular ploy is a call from a company claiming to be "Credit Services" and offering to give you lower credit card rates. They play around with caller ID and have it display anything they want. A common ploy is to display a number similar to yours, but with the last digit different, making them appear local. They often announce, "This is your second and final notice," trying to evoke a sense of urgency.

If you answer, the automated system forwards you to a live person who begins asking a series of "qualifying" questions, such as how much credit card debt you have and what your interest rates are. They then promise that they can help you lower your debt and get you a better rate. Next, they start asking some personal questions. They might ask for the last four digits of your Social Security number so they can "verify your account is in good standing." They might ask you for the code on your credit card "to verify you still have it."

Eventually, they hope to get your credit card number, expiration date, and code so that they can use it to post fraudulent charges. Some people have reported similar callers trying to get their bank information so that they can transfer money out of the accounts.

They hang up right away if you ask them to take you off their list, or stop calling. Similarly, they hang up when they hear words such as *criminal*, *thief*, and other words I'll leave out of this book. Some even reply with insults. They've called me so often, I've played along a few times. I love it when they ask for information on my credit card. I respond by saying, "Can you hold on so I can get it?" I then put the phone in a drawer and go back to work. Once, they stayed on the line for more than three hours waiting for me.

Remember this

Vishing is a form of phishing that uses the phone system or VoIP. Some vishing attempts are fully automated. Others start automated but an attacker takes over at some point during the call.

Privilege Escalation

Privilege escalation occurs when a user or process accesses elevated rights and permissions. When attackers first compromise a system, they often have minimal privileges. However, privilege escalation tactics allow them to get more and more privileges.

For example, imagine hacker Harry is attacking a web server over the Internet. He might only have guest or anonymous access to the system initially, and he can't do much with this access. He uses different techniques during the attack to gain more and more privileges. If he can escalate his privileges high enough, he will have full administrative or root access to the system.

Malware frequently tries to gain access to elevated privileges through the logged-on user. For example, if a user logs on with administrative privileges, the malware can elevate its privileges through the user account.

Many organizations require administrators to have two accounts. They use one account for regular use and one for administrative use. The only time they would log on with the administrator account is when they are performing administrative work. This reduces the time the administrative account is in use, and reduces the potential for privilege escalation if the user's system is infected with malware.

Blocking Malware and Other Attacks

The previous sections described several different methods attackers and criminals use to launch new attacks. However, organizations and individuals can prevent many of these attacks from succeeding with just a few steps. These steps include using anti-malware software and educating users.

Protecting Systems with Anti-Malware Software

Malware is a significant threat for any organization. Administrators commonly implement layered security, or a defense-in-depth plan, to protect against malware. The following bullets list some common security controls used to protect against malware:

- **Anti-malware software on mail servers.** Spam often includes malware as attachments, but anti-malware software can detect and block it. The software strips potentially malicious attachments off the email, and sends a notification to the user explaining what was removed and why.
- **All systems.** All workstations and servers have anti-malware software installed. Servers may have additional, specialized anti-malware software installed depending on the applications running on the servers.
- **Boundaries or firewalls.** Many networks include detection tools that monitor network traffic through the firewall. For example, unified threat management (UTM) inspects network traffic to reduce the risk of malware entering the network. Chapter 3, “Understanding Basic Network Security,” covers UTM systems.

Antivirus Software

Anti-malware software provides protection against many types of malware. You’ll often hear the term *antivirus software* indicating it only protects against viruses. However, the lines have blurred. Viruses aren’t the only threats. Attackers have changed their methodologies using different types of malware, and antivirus software vendors have adapted by including methods to detect and block these new threats. Most antivirus software detects, blocks, and removes several different types of malware, such as viruses, Trojans, worms, rootkits, spyware, and adware.

Antivirus software provides real-time protection and can perform both scheduled and manual scans. The real-time protection continuously monitors the system. For example, when a user visits a web site, antivirus software scans the downloaded web site files and attempts to block malicious code. Similarly, when a user downloads or opens a file, antivirus software scans it before opening it. Scheduled scans occur regularly, such as once a week. If users or technicians detect suspicious activity, they can perform manual scans to check the system.

Antivirus software detects viruses using either signature-based detection or heuristic-based detection.

Signature-Based Detection

Viruses and other malware have known patterns. Signature files (also called data definition files) define the patterns, and the antivirus software scans files for matching patterns. When the software identifies a matching pattern, it reports it as an infection and takes action, such as deleting or quarantining the file.

A quarantined virus is not harmful to the system while it is in quarantine, but it’s still available for analysis. As an example, a security professional could release a quarantined virus into an unprotected but isolated virtual environment for research and study.

Malware developers constantly release new viruses, so it’s important to update signature definition files regularly. Most antivirus software includes the ability to automate the process of checking and downloading updated signature definition files. They typically check for updates several times a day.

It’s also possible to download and install signature files manually. Administrators do this when updating systems that do not have Internet access. When doing so, it’s important for administrators to ensure the signature file has not lost data integrity. They do so by comparing the hash of the signature file posted on the antivirus vendor’s web site with the hash of the downloaded file. To see an example of how to compare hashes, check out the Creating and Comparing Hashes Lab. You can access the online exercises for this book at <http://gcgapremium.com/labs/>.

Remember this

Antivirus software detects and removes malware, such as viruses, Trojans, and worms. Signature-based antivirus software detects known malware based on signature definitions. Heuristic-based software detects previously unknown malware based on behavior.

Heuristic-Based Detection

Some antivirus software includes heuristic-based detection. Heuristic-based detection attempts to detect viruses that were previously unknown and do not have signatures. This includes zero-day exploits, mentioned in Chapter 4, “Securing Your Network.”

Heuristic-based analysis runs questionable code in a sandbox or virtualized environment specifically designed to protect the live environment, while it observes its behavior. Most viruses engage in *viral activities*—actions that can be harmful, but are rarely performed by legitimate programs. The heuristic-based analysis detects these viral activities.

As an example, polymorphic malware adds variations to files when it creates copies. It’s highly unusual for any application to add variations in files like this, and heuristic methods are often successful at detecting polymorphic malware.

Checking File Integrity

Some antivirus scanners use file integrity checkers to detect modified system files. A file integrity checker calculates hashes on system files as a baseline. It then periodically recalculates the hashes on these files and compares them with the hashes in the baseline. If the hashes are ever different, it indicates the system files have been modified. When an antivirus scanner detects a modified file, it sends an alert. Many times these alerts can detect rootkit infections.

When searching for rootkits, antivirus scanners also have the ability to inspect RAM.

Pop-Up Blockers

Most web browsers include a feature to block pop-up ads, and/or support pop-up blocker add-ins that users can configure with the web browser. These block most pop-ups, including the pop-ups described in the “Adware” section earlier in this chapter. Most pop-up blockers allow you to list Uniform Resource Locators (URLs) that allow pop-ups, but block all pop-ups that aren’t on the allowed list. This is similar to creating a whitelist of allowed applications as described in Chapter 5, “Securing Hosts and Data.”

Remember this

The most effective protection against unwanted adware is the use of pop-up blockers in web browsers. Many pop-up blockers support lists of URLs that allow pop-ups.

Spam Filters as Anti-Spam Solutions

Organizations often implement a multipronged approach to block spam. For example, many UTM systems include spam filters to detect and block spam. The output of the UTM goes to an email server. Email servers also have methods of detecting and blocking spam. The email server sends all email to the users, except for what it detects as spam. User systems also have anti-spam filters, or junk mail options, as a final check.

The challenge with any spam filter is to only filter out spam, and never filter out actual email. For example, a company wouldn’t want a spam filter to filter out an email from a customer trying to buy something. Because of this, most spam filters err on the side of allowing spam through rather than potentially marking valid email as spam. Although the science behind spam filtering continues to improve, criminals have also continued to adapt.

Spam filters typically allow you to identify email addresses as safe, or to be blocked. You can add these as individual addresses or entire domains. For example, if you want to ensure you get email from Homer when he sends email from *springfield.com*, you can identify *homer@springfield.com* as a safe email address. If you want to ensure you get all email from *springfield.com*, you can designate *springfield.com* as a safe domain. Similarly, you block either the single email address *homer@springfield.com* or the entire domain *springfield.com*.

Anti-Spyware Software

Anti-spyware software emerged as a separate application that targets spyware. It helps protect a user’s personal information while surfing the Internet. The lines between spyware and malware have become blurry, especially since some spyware has started becoming more malicious. Over time, most antivirus software began including anti-spyware elements. However, not all antivirus software protects against spyware, so some users still run separate anti-spyware applications.

Educating Users

The single best protection against many attacks such as social engineering and other attacks mentioned in this chapter is to train and raise the security awareness of users. Many users simply aren't aware of the attackers' methods. However, once they understand the risks and methods used by social engineers and other attackers, they are less likely to fall prey to these attacks. Similarly, raising users' security awareness helps them recognize and respond appropriately to new threats and security trends.

Security-related awareness and training programs take many forms. Some common methods include formal classes, short informal live training sessions, online courses, posters, newsletters, logon banners, and periodic emails. These programs often keep users aware of new threats and new security trends and alerts, such as new viruses, current phishing attacks, and zero-day exploits.

New Viruses

Criminals are constantly releasing new viruses and some prove to be exceptionally damaging. Many of these require administrators to take quick action to mitigate the threat, but other times, users need to take action.

As an example, security experts discovered Heartbleed in early 2014 and sent out alerts. Administrators managing servers using OpenSSL needed to take steps to eliminate the vulnerability. However, the Heartbleed vulnerability existed on systems for as long as two years and allowed attackers to view information in secure Hypertext Transfer Protocol Secure (HTTPS) sessions. This included passwords and other secure data.

Several alerts told users they should change all the passwords they used on the Internet. The operating system or platform didn't matter. If users logged on to a vulnerable site using a desktop computer, a mobile iOS device, an Android device, or anything else, attackers might have their credentials. Users who were aware of security issues took note and changed their passwords immediately.

If users simply ignore alerts about new threats such as Heartbleed, they remain at risk. Providing training to users on how important these alerts are helps them pay attention and respond when security experts release them.

Phishing Attacks

In addition to releasing new viruses regularly, criminals are also launching new phishing attacks. Some new attempts are tricky and fool many people. The best way to prevent successful attacks is to educate people about what the criminals are doing now.

As an example, criminals hijacked a server in a foreign country and installed drive-by malware on it. If a user visited, the drive-by malware downloaded and installed itself on the user's system. This allowed criminals to access users' systems remotely and gather user data, such as passwords and financial information. Now all they needed to do is get users to visit. They lured users with this message:

From: Eubank Funeral Home

Subject: Death and funeral announcement

For this unprecedented event, we offer our deepest prayers of condolence and invite to you to be present at the celebration of your friends life service on Thursday, February 6, 2014 that will take place at Eubank Funeral Home at 11:00 a.m.

Please find invitation and more detailed information about the farewell ceremony here.

Best wishes and prayers,

Many users clicked and the server in the foreign country downloaded malware onto their systems. Worse, users thought that the Eubank Funeral Home (which is real) launched the attack. The funeral home received 50 to 100 complaints a day from people all over the world. Obviously, this seriously disrupted their business. In time, the Eubank Funeral Home phishing emails stopped. Criminals replaced them with other legitimate funeral homes, such as the Hubbell Funeral Home.

Zero-Day Exploits

Chapter 4 discussed zero-day vulnerabilities and zero-day exploits. As a reminder, a *zero-day vulnerability* is a vulnerability or bug that is unknown to trusted sources, such as operating system and antivirus vendors. Operating system vendors write and release patches once they know about them, but until the vendors know about them, the vulnerability remains. As an example, the Heartbleed vulnerability existed for a couple of years before it was widely published. Up until the time that OpenSSL developers released a fix, everyone using it was vulnerable.

Users might adopt the idea that up-to-date antivirus software will protect them from all malware. This simply isn't true. No matter how great an antivirus company is at identifying new malware, there is always going to be a lag between the time when criminals release the malware and the antivirus company releases new signatures to discover it. This includes malware designed to take advantage of zero-day vulnerabilities.

Remember this

Educating users about new viruses, phishing attacks, and zero-day exploits helps prevent incidents. Zero-day exploits take advantage of vulnerabilities that aren't known by trusted sources, such as operating system vendors and antivirus vendors.

With this in mind, users need to practice safe computing habits. They can't depend on the antivirus software and other technical controls to protect them. Some basic guidelines are:

- Don't click on links within emails from unknown sources (no matter how curious you might be).
- Don't open attachments from unknown sources (malware can be embedded into many different files, such as Portable Document Format (PDF) files, Word documents, Zip files, and more).
- Be wary of free downloads from the Internet (Trojans entice you with something free but include malware).
- Limit information you post on social media sites (criminals use this to answer password reset questions).
- Back up your data regularly (unless you're willing to see it disappear forever).
- Keep your computer up to date with current patches (but beware of zero-day exploits).
- Keep antivirus software up to date (but don't depend on it to catch everything).

Why Social Engineering Works

Social engineers typically use one or more principles, which increase the effectiveness of their attacks. In addition to teaching users about the different social engineering tactics, it's also useful to teach them about these underlying principles. The following sections introduce these topics.

Authority

Many people have grown up to respect authority and are more likely to comply when a person of authority says to do so. As an example, volunteers participating in the Milgram experiment continued to send shocks to unseen subjects even though they could hear them scream in pain, simply because a man in a lab coat told them to continue. They weren't actually sending shocks and the screams were fake, but everything seemed real to the volunteers. Psychologists have repeated these experiments and have seen similar results. Using authority is most effective with impersonation, whaling, and vishing attacks:

- **Impersonation.** Some social engineers impersonate others to get people to do something. For example, many have called users on the phone claiming they work for Microsoft. The Police Virus (mentioned in the "Ransomware" section) attempts to impersonate a law enforcement agency. Other times, social engineers attempt to impersonate a person of authority, such as an executive within a company, or a technician.
- **Whaling.** Executives respect authorities such as legal entities. As an example, the "Whaling" section mentioned how many executives were tricked into opening infected PDF files that looked like official subpoenas.
- **Vishing.** Some attackers use the phone to impersonate authority figures.

Intimidation

In some cases, the attacker attempts to intimidate the victim into taking action. Intimidation might be through bullying tactics, and it is often combined with impersonating someone else. Using intimidation is most effective with impersonation and vishing attacks.

For example, a social engineer might call an executive's receptionist with this request: "Mr. Simpson is about to give a huge presentation to potential customers, but his files are corrupt. He told me to call you and get you to send the files to me immediately so that I can get him set up for his talk." If the receptionist declines, the social engineer can use intimidation tactics by saying something like: "Look, if you want to be responsible for this million-dollar sale falling through, that's fine. I'll tell him you don't want to help."

Note that this tactic can use multiple principles at the same time. In this example, the attacker is combining intimidation with urgency. The receptionist doesn't have much time to respond.

Consensus/Social Proof

People are often more willing to like something that other people like. Some attackers take advantage of this by creating web sites with fake testimonials that promote a product. For example, criminals have set up some web sites with dozens of testimonials listing all the benefits of their fake antivirus software (rogueware). If users search the Internet before downloading the rogueware, they will come across these web sites, and might believe that other real people are vouching for the product.

Using consensus/social proof is most effective with Trojans and hoaxes. Victims are more likely to install a Trojan if everyone seems to indicate it's safe. Similarly, if a person suspects a virus notice is a just a hoax, but everyone seems to be saying it's real, the victim is more likely to be tricked.

Scarcity

People are often encouraged to take action when they think there is a limited quantity. As an example of scarcity, think of Apple iPhones. When Apple first releases the new version, they typically sell out quickly. A phishing email can take advantage of this and encourage users to click a link for exclusive access to a new product. If the users click, they'll end up at a malicious web site. Scarcity is often effective with phishing and Trojan attacks. People make quick decisions without thinking them through.

Urgency

Some attacks use urgency as a technique to encourage people to take action now. As an example, the CryptoLocker ransomware virus mentioned in this chapter uses the scarcity principle with a countdown timer. Victims have 72 hours before they'll lose all their data, and each time they look at their computer, they'll see the timer counting down.

Using urgency is most effective with ransomware, phishing, vishing, whaling, and hoaxes. For example, phishing emails with malicious links might indicate that there are a limited number of products at a certain price, so the user should "Click Now." Executives might be tricked into thinking a subpoena requires immediate action. Many virus hoaxes have a deadline such as at 4:00 p.m. when the hoax claims the virus will cause the damage.

Remember this

Many of the reasons that social engineers are effective are because they use psychology-based techniques to overcome users' objectives. Scarcity and urgency are two techniques that encourage immediate action.

Familiarity/Liking

If you like someone, you are more likely to do what the person asks. This is why so many big companies hire well-liked celebrities. And, it's also why they fire them when those celebrities become embroiled in a scandal that affects their credibility.

Some social engineers attempt to build rapport with the victim to build a relationship before launching the attack. This principle is most effective with shoulder surfing and tailgating attacks:

- **Shoulder surfing.** People are more likely to accept someone looking over their shoulder when they are familiar with the other person, or they like them. In contrast, if people don't know or don't like someone, they are more likely to recognize a shoulder surfing attack and stop it immediately.
- **Tailgating.** People are much more likely to allow someone to tailgate behind them if they know the person or like the person. Some social engineers use a simple, disarming smile to get the other person to like them.

Trust

In addition to familiarity/liking, some social engineers attempt to build a trusting relationship between them and the victim. This often takes a little time, but the reward for the criminal can be worth it. Vishing attacks often use this method.

As an example, someone identifying himself as a security expert once called me. He said he was working for some company with "Secure" in its name, and they noticed that my computer was sending out errors. He stressed a couple of times that they deploy and support Windows systems. The company name and their experience was an attempt to start building trust.

He then guided me through the process of opening Event Viewer and viewing some errors on my system. He asked me to describe what I saw and eventually said, "Oh my God!" with the voice of a well-seasoned actor. He explained that this indicated my computer was seriously infected. In reality, the errors were trivial.

After seriously explaining how much trouble I was in with my computer, he then added a smile to his voice and said, "But this is your lucky day. I'm going to help you." He offered to guide me through the process of fixing my computer before the malware damaged it permanently.

All of this was to build trust. At this point, he went in for the kill. He had me open up the Run window and type in a web site address and asked me to click OK. This is where I stopped. I didn't click OK. I tried to get him to answer some questions but he was evasive. Eventually, I heard a click. My "lucky day" experience with this social engineering criminal was over.

The link probably would have taken me to a malicious web site ready with a drive-by download. Possibly the attacker was going to guide me through the process of installing rogueware on my system. If my system objected with an error, I'm betting he would have been ready with a soothing voice saying, "That's normal. Just click OK. Trust me." He spent a lot of time with me. I suspect that they've been quite successful with this ruse with many other people.

Chapter 6 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Understanding Malware Types

- Malware includes several different types of malicious code, including viruses, worms, ransomware, logic bombs, rootkits, backdoors, and more.
- An armored virus uses one or more techniques to make it difficult to reverse engineer. Polymorphic malware changes to make it difficult to detect.
- A worm is self-replicating, unlike a virus, which must be executed.
- A logic bomb executes in response to an event, such as a day, time, or condition. Malicious insiders have planted logic bombs into existing systems, and these logic bombs have delivered their payload after the employee left the company.
- Backdoors provide another way of accessing a system. Malware often inserts backdoors into systems, giving attackers remote access to systems.
- A Trojan appears to be one thing, such as pirated software or free antivirus software, but is something malicious. Drive-by downloads attempt to infect systems with Trojans.
- A botnet is group of computers called zombies controlled through a command-and-control server. Attackers use malware to join computers to botnets.
- Ransomware is a type of malware that takes control of a user's system or data. Criminals attempt to extort payment as ransom combined with threats of damaging a user's system or data if the user doesn't pay.
- Rootkits take root-level or kernel-level control of a system. They hide their processes to avoid detection. They can remove user privileges and modify system files.
- Adware often causes pop-up windows to appear with advertisements.
- Spyware is software installed on user systems without the user's knowledge or consent and it monitors the user's activities. It can result in the loss of confidentiality as it steals user secrets.

Recognizing Common Attacks

- Social engineering is the practice of using social tactics to gain information or trick users into performing an action they wouldn't normally take.
- Social engineering attacks can occur in person, over the phone, while surfing the Internet, and via email. Many social engineers attempt to impersonate others.
- Shoulder surfing is an attempt to gain unauthorized information through casual observation, such as looking over someone's shoulder, or monitoring screens with a camera.
- Tailgating is the practice of one person following closely behind another without showing credentials. Mantraps help prevent tailgating. Cameras with recording capabilities are a cheaper substitute to deter tailgating.
- Dumpster divers search through trash looking for information. Shredding or burning documents reduces the risk of dumpster diving.
- Spam is unwanted or unsolicited email. Attackers often use spam in different types of attacks.
- Phishing is the practice of sending email to users with the purpose of tricking them into revealing sensitive information, installing malware, or clicking on a link.
- Spear phishing and whaling are types of phishing. Spear phishing targets specific groups of users and whaling targets high-level executives.
- Vishing is a form of phishing that uses voice over the telephone and often uses Voice over IP (VoIP). Some vishing attacks start with a recorded voice and then switch over to a live person.

Blocking Malware and Other Attacks

- Antivirus software can detect and block different types of malware, such as worms, viruses, and Trojans. Antivirus software uses signatures to detect known malware.
- When downloading signatures manually, hashes can verify the integrity of signature files.
- Antivirus software typically includes a file integrity checker to detect files modified by a rootkit.
- Pop-up blockers can block many pop-up windows used by adware.
- Anti-spam software attempts to block unsolicited email. You can configure a spam filter to block individual email addresses and email domains.
- Anti-spyware software helps protect users' personal information while online by detecting and blocking spyware. Some antivirus software

applications include anti-spyware elements.

- Security-related awareness and training programs help users learn about new threats and security trends, such as new viruses, new phishing attacks, and zero-day exploits. Zero-day exploits take advantage of vulnerabilities that aren't known by trusted sources.
- Social engineers and other criminals employ several principles to help increase the effectiveness of their attacks. They are authority, intimidation, consensus/social proof, scarcity, urgency, familiarity/liking, and trust.

Chapter 6 Practice Questions

1. Which of the following types of malware is the MOST difficult to reverse engineer?
 - A. Logic bomb
 - B. Trojan
 - C. Armored virus
 - D. Ransomware
2. Recently, malware on a company computer destroyed several important files after it detected that Homer was no longer employed at the company. Which of the following BEST identifies this malware?
 - A. Logic bomb
 - B. Rootkit
 - C. Backdoor
 - D. Adware
3. A recent antivirus scan on a server detected a Trojan. A technician removed the Trojan, but a security administrator expressed concern that unauthorized personnel might be able to access data on the server. The security administrator decided to check the server further. Of the following choices, what is the administrator MOST likely looking for on this server?
 - A. Backdoor
 - B. Logic bomb
 - C. Rootkit
 - D. Botnet
4. After Maggie turned on her computer, she saw a message indicating that unless she made a payment, her hard drive would be formatted. What does this indicate?
 - A. Armored virus
 - B. Ransomware
 - C. Backdoor
 - D. Trojan
5. A security administrator recently noticed abnormal activity on a workstation. It is connecting to computers outside the organization's internal network, using uncommon ports. Using a security toolkit, the administrator discovered the computer is also running several hidden processes. Which of the following choices BEST indicates what the administrator has found?
 - A. Rootkit
 - B. Backdoor
 - C. Spam
 - D. Trojan
6. What type of malware uses marketing pop-ups and does not attempt to hide itself?
 - A. Blocker
 - B. Rootkit
 - C. Trojans
 - D. Adware
7. Of the following malware types, which one is MOST likely to monitor a user's computer?
 - A. Trojan

- B. Spyware
- C. Adware
- D. Ransomware

8. Lisa is a database administrator and received a phone call from someone identifying himself as a technician working with a known hardware vendor. The technician said he's aware of a problem with database servers they've sold, but it only affects certain operating system versions. He asks Lisa what operating system the company is running on its database servers. Which of the following choices is the BEST response from Lisa?

- A. Let the caller know what operating system and versions are running on the database servers to determine if any further action is needed.
- B. Thank the caller and end the call, report the call to her supervisor, and independently check the vendor for issues.
- C. Ask the caller for his phone number so that she can call him back after checking the servers.
- D. Contact law enforcement personnel.

9. A security administrator at a shopping mall discovered two wireless cameras pointing at an automatic teller machine. These cameras were not installed by mall personnel and are not authorized. What is the MOST likely goal of these cameras?

- A. Tailgating
- B. Dumpster diving
- C. Vishing
- D. Shoulder surfing

10. Bart is in a break area outside the office. He told Lisa that he forgot his badge inside and asked Lisa to let him follow her when she goes back inside.

What does this describe?

- A. Spear phishing
- B. Whaling
- C. Mantrap
- D. Tailgating

11. An organization's security policy requires employees to incinerate paper documents. Of the following choices, which type of attack is this MOST likely to prevent?

- A. Shoulder surfing
- B. Tailgating
- C. Vishing
- D. Dumpster diving

12. While cleaning out his desk, Bart threw several papers containing PII into the recycle bin. Which type of attack can exploit this action?

- A. Vishing
- B. Dumpster diving
- C. Shoulder surfing
- D. Tailgating

13. Marge reports that she keeps receiving unwanted emails about personal loans. What does this describe?

- A. Phishing
- B. Spear phishing
- C. Spam
- D. Vishing

14. A recent spear phishing attack that appeared to come from your organization's CEO resulted in several employees revealing their passwords to attackers. Management wants to implement a security control to provide assurances to employees that email that appears to come from the CEO actually came from the CEO. Which of the following should be implemented?

- A. Digital signatures
- B. Spam filter
- C. Training
- D. Metrics

15. Attackers are targeting C-level executives in your organization. Which type of attack is this?
- A. Phishing
 - B. Vishing
 - C. Spam
 - D. Whaling
16. You manage a group of computers in an isolated network without Internet access. You need to update the antivirus definitions manually on these computers. Which of the following choices is the MOST important concern?
- A. Running a full scan of the systems before installing the new definitions
 - B. Running a full scan of the systems after installing the new definitions
 - C. Ensuring the definition file hash is equal to the hash on the antivirus vendor’s web site
 - D. Ensuring the update includes all signature definitions
17. A user wants to reduce the threat of an attacker capturing her personal information while she surfs the Internet. Which of the following is the BEST choice?
- A. Antivirus software
 - B. Anti-spyware software
 - C. Pop-up blocker
 - D. Whitelisting
18. Bart is complaining that new browser windows keep opening on his computer. Which of the following is the BEST choice to stop these in the future?
- A. Malware
 - B. Adware
 - C. Pop-up blocker
 - D. Antivirus software
19. Your organization recently suffered a loss from malware that wasn’t previously known by any trusted sources. Which type of attack is this?
- A. Phishing attack
 - B. Zero-day
 - C. Buffer overflow
 - D. Integer overflow
20. Homer received an email advertising the newest version of a popular smartphone, which is not available elsewhere. It includes a malicious link. Which of the following principles is the email author using?
- A. Authority
 - B. Intimidation
 - C. Scarcity
 - D. Trust

Chapter 6 Practice Question Answers

1. **C.** An armored virus uses one or more techniques to make it difficult for antivirus researchers to reverse engineer it. A logic bomb executes in response to an event, but it is often implemented with simple code. A Trojan appears to be something beneficial, but it includes a malicious component. Ransomware takes control of a user’s system or data and then demands payment as ransom.
2. **A.** A logic bomb executes in response to an event. In this scenario, the logic bomb is delivering its payload when it detects that Homer is no longer employed at the company. A rootkit doesn’t respond to an event. A backdoor provides another method of accessing a system, but it does not delete files. Adware uses advertising methods, such as pop-up windows.
3. **A.** The security administrator is most likely looking for a backdoor because Trojans commonly create backdoors, and a backdoor allows unauthorized personnel to access data on the system. Logic bombs and rootkits can create backdoor accounts, but Trojans don’t create logic bombs and would rarely install a rootkit. The computer might be joined to a botnet, but it wouldn’t be a botnet.
4. **B.** Ransomware attempts to take control of a user’s system or data and then demands ransom to return control. An armored virus uses one or more

- techniques to make it more difficult to reverse engineer. It's possible that Maggie's computer was infected with a Trojan, which created a backdoor. However, not all Trojans or backdoor accounts demand payment as ransom.
5. **A.** A rootkit typically runs processes that are hidden and it also attempts to connect to computers via the Internet. Although an attacker might have used a backdoor to gain access to the user's computer and install the rootkit, backdoors don't run hidden processes. Spam is unwanted email and is unrelated to this question. A Trojan is malware that looks like it's beneficial, but is malicious.
6. **D.** Adware commonly causes pop-up windows to appear with marketing advertisements and adware doesn't try to hide itself. Many web browsers include pop-up blockers that block these pop-ups. A rootkit does attempt to hide itself and keep any rootkit processes hidden. Trojans perform some malicious activity such as creating a backdoor account, and they hide their activity.
7. **B.** Spyware monitors a user's computer and activity. Trojans often install backdoor accounts, but they don't necessarily monitor systems and activity. Adware typically causes pop-up windows for advertising, and although it might monitor the user to target ads, not all adware monitors users. Ransomware is primarily concerned with getting the user to make a ransom payment.
8. **B.** This sounds like a social engineering attack where the caller is attempting to get information on the servers, so it's appropriate to end the call, report the call to a supervisor, and independently check the vendor for potential issues. It is not appropriate to give external personnel information on internal systems from a single phone call. The caller has not committed a crime by asking questions, so it is not appropriate to contact law enforcement personnel.
9. **D.** Shoulder surfing is the practice of peering over a person's shoulder to discover information. In this scenario, the attacker is using the wireless cameras to discover PINs as users enter them. Tailgating is the practice of following closely behind someone else without using credentials. Dumpster diving is the practice of searching trash dumpsters for information. Vishing is a form of phishing using the phone.
10. **D.** Tailgating is the practice of following closely behind someone else without using credentials. In this scenario, Bart might be an employee who forgot his badge, or he might be a social engineer trying to get in by tailgating. Mantraps prevent tailgating. Spear phishing and whaling are two types of phishing with email.
11. **D.** Dumpster diving is the practice of looking for documents in the trash dumpsters, but shredding or incinerating documents ensures dumpster divers cannot retrieve any paper documents. Shoulder surfers attempt to view something on a monitor or other screen, not papers. Tailgating refers to entering a secure area by following someone else. Vishing is a form of phishing using the phone.
12. **B.** Dumpster divers look through trash or recycling containers for valuable paperwork, such as documents that include Personally Identifiable Information (PII). Instead, paperwork should be shredded or incinerated. Vishing is a form of phishing that uses the phone. Shoulder surfers attempt to view monitors or screens, not papers. Tailgating is the practice of following closely behind someone else, without using proper credentials.
13. **C.** Spam is unwanted emails from any source. Phishing and spear phishing are types of attacks using email. Vishing is similar to phishing but it uses telephone technology.
14. **A.** A digital signature provides assurances of who sent an email and meets the goal of this scenario. Although a spam filter might filter a spear phishing attack, it does not provide assurances about who sent an email. A training program would help educate employees about attacks and would help prevent the success of these attacks, but it doesn't provide assurances about who sent an email. Metrics can measure the success of a training program.
15. **D.** Whaling is a type of phishing that targets high-level executives, such as CEOs, CIOs, and CFOs. Because whaling is more specific than phishing, phishing isn't the best answer. Vishing is similar to phishing, but it uses the phone instead. Spam is unwanted email, but spam isn't necessarily malicious.
16. **C.** When downloading files as important as antivirus definitions, it's important to ensure they do not lose data integrity, and you can do so by verifying the hashes. It's not necessary to run a full scan either before or after installing new definitions, but the new definitions will help.
17. **B.** Anti-spyware is the best choice to protect an individual's personal information while online. Many antivirus software applications include anti-spyware components, but not all of them do. A pop-up blocker prevents pop-up windows, caused by adware. Whitelisting identifies specific applications authorized on a system, but does not necessarily prevent the theft of personal information.
18. **C.** A pop-up blocker is the best choice to stop these windows, which are commonly called pop-up windows. They might be the result of malware or adware, but more malware or adware will not stop them. Some antivirus software may block the pop-ups, but a pop-up blocker is the best choice.
19. **B.** A zero-day exploit is one that isn't known by trusted sources such as antivirus vendors or operating system vendors. Trusted sources know about many phishing attacks, buffer overflow attacks, and integer overflow attacks.
20. **C.** The attacker is using scarcity to entice the user to click the link. A user might realize that clicking on links from unknown sources is risky, but the temptation of getting the new smartphone might cause the user to ignore the risk.

Chapter 7

Identifying Advanced Attacks

CompTIA Security+ objectives covered in this chapter:

1.2 Given a scenario, use secure network administration principles.

- Flood guards

3.2 Summarize various types of attacks.

- Man-in-the-middle, DDoS, DoS, Replay, Smurf attack, Spoofing, Xmas attack, Pharming, DNS poisoning and ARP poisoning, Transitive access, Client-side attacks
- Password attacks (Brute force, Dictionary attacks, Hybrid, Birthday attacks, Rainbow tables)
- Typo squatting/URL hijacking, Watering hole attack

3.4 Explain types of wireless attacks.

- Replay attacks

3.5 Explain types of application attacks.

- Cross-site scripting, SQL injection, LDAP injection, XML injection, Directory traversal/command injection, Buffer overflow, Integer overflow, Zero-day, Cookies and attachments, LSO (Locally Shared Objects), Flash Cookies, Malicious add-ons, Session hijacking, Header manipulation, Arbitrary code execution / remote code execution

4.1 Explain the importance of application security controls and techniques.

- Fuzzing
- Secure coding concepts (Error and exception handling, Input validation)
- Cross-site scripting prevention, Cross-site Request Forgery (XSRF) prevention, NoSQL databases vs. SQL databases, Server-side vs. Client-side validation

If there's one thing that's abundant in the IT world, it is attacks and attackers. Attackers lurk almost everywhere. If you have computer systems, you can't escape them. However, you can be proactive in identifying the different types of attacks and take steps to prevent them, or at least prevent their effectiveness. This chapter covers a wide assortment of attacks from different sources and provides some insight into preventing many of them.

Comparing Common Attacks

This section summarizes several common types of attacks launched against systems and networks. Some of them are generic, such as denial-of-service attacks, and others are very specific. It's important to realize that effective countermeasures exist for all of the attacks listed in this book. However, attackers are actively working on beating the countermeasures. As they do, security professionals create additional countermeasures and the attackers try to beat them. The battle continues daily.

The goal in this section is to become aware of many of the well-known attacks. By understanding these, you'll be better prepared to comprehend the improved attacks as they emerge and the improved countermeasures.

Spoofing

Spoofing occurs when one person or entity impersonates or masquerades as someone or something else. Many different types of attacks use spoofing.

For example, Chapter 4, “Securing Your Network,” mentioned how wireless attackers can bypass media access control (MAC) address filtering by spoofing the MAC address of authorized systems. The smurf attack (mentioned later in this chapter) spoofs the source IP address by replacing the original address with the IP address of the victim.

Email spoofing occurs when someone changes the “From” address in an email to make it appear as though the email is coming from someone else. The email may come from an attacker, but the spoofed address attempts to hide the attacker’s identity and make the email look valid. This is actually easy to do with most email software.

DoS Versus DDoS

A *denial-of-service (DoS) attack* is an attack from one attacker against one target. A *distributed denial-of-service (DDoS) attack* is an attack from two or more computers against a single target. DDoS attacks often include sustained, abnormally high network traffic on the network interface card of the attacked computer. Other system resource usage (such as the processor and memory usage) will also be abnormally high. The goal of both is to prevent legitimate users from accessing services on the target computer.

Many DoS and DDoS attacks attempt to consume resources on the target computer. For example, a SYN (synchronize) flood attack consumes memory resources by flooding a system with half-open connections.

Remember this

A denial-of-service (DoS) attack is an attack from a single source that attempts to disrupt the services provided by another system. A distributed denial-of-service (DDoS) attack includes multiple computers attacking a single target. DDoS attacks typically include sustained, abnormally high network traffic.

Smurf Attacks

A *smurf attack* spoofs the source address of a directed broadcast ping packet to flood a victim with ping replies. It's worthwhile to break this down:

- **A ping is normally unicast—one computer to one computer.** A ping sends ICMP echo requests to one computer, and the receiving computer responds with ICMP echo responses.
- **The smurf attack sends the ping out as a broadcast.** In a broadcast, one computer sends the packet to all other computers in the subnet.
- **The smurf attack spoofs the source IP.** If the source IP address isn't changed, the computer sending out the broadcast ping will get flooded with the ICMP replies. Instead, the smurf attack substitutes the source IP with the IP address of the victim, and the victim gets flooded with these ICMP replies.

Smurf attacks typically use directed broadcasts with an amplifying network similar to what's shown in Figure 7.1. The attacker sends a directed broadcast ping through a router into another network, spoofing the source IP address with the victim's IP address of 10.80.5.1 instead. Each of the computers in the amplifying network responds by flooding the victim with ping responses. Any network can become an amplifying network.




Figure 7.1: Smurf attack through an amplifying network

Most routers block directed broadcasts by default. This is especially important for any border routers between a public network such as the Internet and a private network. In Figure 7.1, Router 1 is a border router. Blocking directed broadcasts prevents an internal network from becoming part of an attack as an amplifying network. For example, if the attacker spoofs an IP address of a computer on the Internet, computers in the amplifying network will attack that external computer.

Remember this

Smurf attacks typically use directed broadcasts to launch attacks through amplifying networks. Disabling directed broadcasts on routers mitigates the threat. It's especially important to ensure directed broadcasts are disabled on routers bordering on the Internet to ensure internal networks are not used as amplifying networks.

SYN Flood Attacks

The SYN flood attack is a common attack used against servers on the Internet. They are easy for attackers to launch, difficult to stop, and can cause significant problems. The SYN flood attack disrupts the TCP handshake process and can prevent legitimate clients from connecting.

Chapter 3, “Understanding Basic Network Security,” explained how TCP sessions use a three-way handshake when establishing a session. As a reminder, two systems normally start a TCP session by exchanging three packets in a TCP handshake. For example, when a client establishes a session with a server, it takes the following steps:

1. The client sends a SYN (synchronize) packet to the server.
2. The server responds with a SYN/ACK (synchronize/acknowledge) packet.
3. The client completes the handshake by sending an ACK (acknowledge) packet. After establishing the session, the two systems exchange data.

However, in a SYN flood attack, the attacker never completes the handshake by sending the ACK packet. Additionally, the attacker sends a barrage of SYN packets, leaving the server with multiple half-open connections. Figure 7.2 compares a normal TCP handshake with the start of a SYN flood attack.



Figure 7.2: TCP handshake and SYN flood attack

In some cases, these half-open connections can consume a server’s resources while it is waiting for the third packet, and it can actually crash. More often though, the server limits the number of these half-open connections. Once the limit is reached, the server won’t accept any new connections, blocking connections from legitimate users. For example, Linux systems support an iptables command that can set a threshold for SYN packets, blocking them after the threshold is set. Although this prevents the SYN flood attack from crashing the system, it also denies service to legitimate clients.

Flood Guards

Flood guards use a variety of different methods to protect against SYN flood attacks. Many firewalls and intrusion detection systems include flood guards, which are simply techniques to limit the success of a SYN flood attack. Additionally, some vendors sell flood guard appliances dedicated to detecting and blocking these attacks.

One method of detecting and blocking these attacks is by identifying the source IP address. If a single source IP address is initiating these half-open connections, but never completing them, a flood guard can block all traffic from this IP. However, attackers now commonly spoof the source IP address in each SYN flood packet. Additionally, attackers often launch attacks from multiple systems at the same time, making it difficult to identify legitimate traffic from attacking traffic.

Another method is to dynamically adjust the time a system waits for the third packet. For example, the system may normally wait 75 seconds for an ACK after sending the SYN/ACK packet. After sensing a barrage of SYN packets, it can reduce the time it waits for the ACK.

There is a lot more depth to SYN flood attacks and methods used to mitigate them. Additionally, attacks and mitigation techniques continue to evolve. If you’re interested in digging deeper, check out RFC 4987, “TCP SYN Flooding Attacks and Common Mitigations,” here: <http://tools.ietf.org/html/rfc4987>.

Remember this

A SYN flood attack disrupts the TCP initiation process by withholding the third packet of the TCP three-way handshake. Flood guards protect against SYN flood attacks.

Xmas Attacks

The *Xmas attack*, also called a Christmas tree attack, is a type of port scan used to identify underlying details of an operating system. For example, it can help determine if the scanned system is running a Microsoft-based operating system or a Linux-based operating system.

A typical port scan (introduced in Chapter 3) attempts to learn what ports are open on a system. Based on what ports are open, the port scanner can detect what services and protocols are running on a system. For example, if port 80 is open, it's very likely that the Hypertext Transfer Protocol (HTTP) protocol is running on the system because the well-known port for HTTP is port 80.

However, the Xmas attack goes farther than a typical port scan. It has several bits set in the packet header and is reminiscent of lights lit in a Christmas tree. As least *someone* thought it looked like a Christmas tree and decided to name it a Christmas tree attack, or Xmas attack.

More importantly, the Xmas attack sets specific flags within the TCP packet header. Different operating systems respond to these flags in specific ways. Attackers can analyze the response and determine the operating system of the remote system in addition to what ports are open. In many cases, the attacker can even determine the version of the responding system.

The Xmas attack is often used as reconnaissance in an overall attack. It doesn't cause damage itself. However, attackers use the information they gain from the Xmas attack to launch other attacks. Most intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) can detect these attacks.

Man-in-the-Middle Attacks

A *man-in-the-middle (MITM) attack* is a form of active interception or active eavesdropping. It uses a separate computer that accepts traffic from each party in a conversation and forwards the traffic between the two. The two computers are unaware of the MITM computer, and it can interrupt the traffic at will or insert malicious code.

For example, imagine that Maggie and Bart are exchanging information with their two computers over a network. If hacker Harry can launch an MITM attack from a third computer, he will be able to intercept all traffic. Maggie and Bart still receive all the information, so they are unaware of the attack. However, hacker Harry also receives all the information. Because the MITM computer can control the entire conversation, it is easy to insert malicious code and send it to the computers. The “ARP Poisoning” section later in this chapter shows how ARP poisoning can be used to launch an MITM attack.

Kerberos helps prevent man-in-the-middle attacks with mutual authentication. It doesn’t allow a malicious system to insert itself in the middle of the conversation without the knowledge of the other two systems.

Replay Attacks

A *replay attack* is one where an attacker replays data that was already part of a communication session. In a replay attack, a third party attempts to impersonate a client that is involved in the original session. Replay attacks can occur on both wired and wireless networks.

As an example, Maggie and Bart may initiate a session with each other. During the communication, each client authenticates with the other by passing authentication credentials to the other system. Hacker Harry intercepts all the data, including the credentials, and later initiates a conversation with Maggie pretending to be Bart. When Maggie challenges hacker Harry, he sends Bart’s credentials.

Many protocols use timestamps and sequence numbers to thwart replay attacks. For example, Kerberos, covered in Chapter 1, “Mastering Security Basics,” helps prevent replay attacks with timestamped tickets.

Remember this

Replay attacks capture data in a session with the intent of later impersonating one of the parties in the session. Timestamps and sequence numbers are effective countermeasures against replay attacks.

Password Attacks

Password attacks attempt to discover or bypass passwords used for authentication on systems and networks, and for different types of files.

Although there are many attack methods, they fall into two generic categories: online password attacks and offline password attacks.

An *online password attack* attempts to discover a password from an online system. For example, an attacker trying to log on to an account by trying to guess a user's password is an online attack.

Offline password attacks attempt to discover passwords from a captured database or captured packet scan. For example, when attackers hack into a web site causing a data breach, they can download entire databases. They then perform offline attacks to discover the passwords contained within the databases. Similarly, Wi-Fi Protected Access (WPA) cracking attacks (discussed in Chapter 4) first capture the four-way handshake when WPA wireless clients authenticate with a wireless access point (WAP). After capturing these packets, they perform an offline attack to discover the WPA password.

The following sections cover some specific types of password attacks.

Brute Force Attacks

A *brute force attack* attempts to guess all possible character combinations. One of the best protections against offline brute force attacks is to use complex passwords, as described in Chapter 1. Complex passwords include a mix of uppercase letters, lowercase letters, numbers, and special characters. Additionally, longer passwords are much more difficult to crack than shorter passwords.

Account lockout policies (also covered in Chapter 1) are effective against online brute force attacks. An account lockout setting locks an account after the user enters the incorrect password a preset number of times.

Remember this

Account lockout policies protect against online brute force password attacks. Complex passwords of sufficient length protect against offline brute force attacks.

Dictionary Attacks

One of the original password attacks uses a dictionary of words and just attempts to use every word in the dictionary to see if it works. Dictionaries used in these attacks evolved over time and included many of the common passwords that uneducated users configured for their accounts. For example, even though 12345 isn't a dictionary word, many people use it as a password, so characters such as these have been added to many dictionaries used by dictionary attack tools.

These attacks are thwarted by using complex passwords. A complex password will not include words in a dictionary.

Password Hashes

Most systems don't store the actual password for an account. Instead, they store a hash of the password. Hash attacks attack the hash of a password instead of the password. Chapter 1 introduced hashing and, as a reminder, a *hash* is simply a number created with a hashing algorithm such as MD5 or SHA-1. A system can use a hashing algorithm such as Message Digest 5 (MD5) to create a hash of a password.

As an example, if a user's password is IC@nP@\$\$S3curity+, the system calculates the hash and stores it instead. In this example, the MD5 hash is 75c8ac11c86ca966b58166187589cc15. Later, a user authenticates with a username and password. The system then calculates the hash of the password entered by the user and compares the calculated hash against the stored hash. If they are correct, it indicates the user entered the correct password.

Similarly, systems rarely send passwords across a network. Instead, they send the hash of the password and normally in an encrypted format.

Unfortunately, tools are available to discover many hashed passwords. For example, MD5 Online (<http://www.md5online.org/>) allows you to enter a hash and it gives you the text of the password. If the password is 12345, the hash is 827ccb0eea8a706c4c34a16891f84e7b. If you enter that hash into MD5 Online, it returns the password of 12345 in about a second. MD5 Online uses a database of hashed words from a dictionary. If the hash matches a database entry, the site returns the password.

Birthday Attacks

A *birthday attack* is named after the birthday paradox in mathematical probability theory. The birthday paradox states that for any random group of 23 people, there is a 50 percent chance that 2 of them have the same birthday. This is not the same year, but instead one of the 365 days in any year.

In a birthday attack, an attacker is able to create a password that produces the same hash as the user's actual password. This is also known as a hash collision.

A *hash collision* occurs when the hashing algorithm creates the same hash from different passwords. This is not desirable. As an example, imagine a simple hashing algorithm creates three-digit hashes. The password "success" might create a hash of 123 and the password "passed" might create the same hash of 123. In this scenario, an attacker could use either "success" or "passed" as the password and both would work.

Birthday attacks on hashes are thwarted by increasing the number of bits used in the hash to increase the number of possible hashes. For example, the MD5 algorithm uses 128 bits and is susceptible to birthday attacks. Secure Hash Algorithm version 2 (SHA-2) can use as many as 512 bits and it is not susceptible to birthday attacks.

Rainbow Table Attacks

Rainbow table attacks are a type of attack that attempts to discover the password from the hash. However, they use rainbow tables, which are huge databases of precomputed hashes. It helps to look at the process of how some password crackers discover passwords without a rainbow table. Assume that an attacker has the hash of a password. The application can use the following steps to crack it:

1. The application guesses a password (or uses a password from a dictionary).
2. The application hashes the guessed password.
3. The application compares the original password hash with the guessed password hash. If they are the same, the application knows the password.
4. If they aren't the same, the application repeats steps 1 through 3 until finding a match.

From a computing perspective, the most time-consuming part of these steps is hashing the guessed password in step 2. However, by using rainbow tables, applications eliminate this step. Rainbow tables are huge databases storing passwords and their calculated hashes. Some rainbow tables are as large as 160 GB in size, and they include hashes for every possible combination of characters up to eight characters in length. Larger rainbow tables are also available.

In a rainbow table attack, the application simply compares the hash of the original password against hashes stored in the rainbow table. When the application finds a match, it identifies the password used to create the hash (or at least text that can reproduce the hash of the original password). Admittedly, this is a simplistic explanation of a rainbow table attack, but it is adequate if you don't plan on writing the algorithm to create your own rainbow table attack software.

Salting passwords is a common method of preventing rainbow table attacks, along with other password attacks such as dictionary attacks. A *salt* is a set of random data such as two additional characters. Password salting adds these additional characters to a password before hashing it. These additional characters add complexity to the password, and also result in a different hash than the system would create using the original password. This causes password attacks that compare hashes to fail.

Chapter 10, “Understanding Cryptography,” covers bcrypt and Password-Based Key Derivation Function 2 (PBKDF2). Both use salting techniques to increase the complexity of passwords and thwart brute force and rainbow table attacks.

Remember this

Passwords are typically stored as hashes. Salting adds random text to passwords before hashing them and thwarts many password attacks.

Hybrid Attacks

A *hybrid attack* uses a combination of two or more attacks to crack a password. As an example, a dictionary attack can use a dictionary of words, but also combine it with a brute force attack by modifying the words. For example, after using all the words in the dictionary, a password cracker can append all the words with a number such as 1 and try them.

DNS Attacks

Chapter 3 covers Domain Name System (DNS) in much more depth, but as a reminder, DNS resolves host names to IP addresses. This eliminates the need for you and me to have to remember the IP address for web sites. Instead, we simply type the name into the browser, and it connects. For example, if you type in *gcgapremium.com* as the Uniform Resource Locator (URL) in your web browser, your system queries a DNS server for the IP address. DNS responds with the correct IP address and your system connects to the web site using the IP address.

DNS also provides reverse lookups. In a reverse lookup, a client sends an IP address to a DNS server with a request to resolve it to a name. Some applications use this as a rudimentary security mechanism to detect spoofing. For example, an attacker may try to spoof the computer's identity by using a different name during a session. However, the Transmission Control Protocol/Internet Protocol (TCP/IP) packets in the session include the IP address of the masquerading system and a reverse lookup shows the system's actual name. If the names are different, it shows suspicious activity. Reverse lookups are not 100 percent reliable because reverse lookup records are optional on DNS servers. However, they are useful when they're available.

Two attacks against DNS services are DNS poisoning and pharming.

DNS Poisoning Attacks

A *DNS poisoning attack* attempts to modify or corrupt DNS results. For example, a successful DNS poisoning attack can modify the IP address associated with *google.com* and replace it with the IP address to a malicious web site. Each time a user queries DNS for the IP address of *google.com*, the DNS server responds with the IP address of the malicious web site.

There have been several successful DNS poisoning attacks over the years. Many current DNS servers use Domain Name System Security Extensions (DNSSEC) to protect the DNS records and prevent DNS poisoning attacks.

Pharming Attacks

A *pharming attack* is another type of attack that manipulates the DNS name resolution process. It either tries to corrupt the DNS server or the DNS client. Just as a DNS poisoning attack can redirect users to different web sites, a successful pharming attack redirects a user to a different web site.

Pharming attacks on the client computer modify the hosts file used on Windows systems. This file is located in the *C:\Windows\System32\drivers\etc* folder. A default entry in the hosts file resolves the host name to the IP address of 127.0.0.1. If an attacker is able to modify other entries, he can cause systems to use that IP address instead of querying DNS. Many viruses have done this in the past. Here's an example of a corrupted hosts file that modifies the entry for google.com:

```
127.0.0.1      localhost
204.79.197.200 google.com
```

Microsoft Bing uses the IP address of 204.79.197.200, so this modified hosts file will now cause all attempts to use Google to be redirected to Bing instead. Practical jokers might do this to a friend's computer and it isn't malicious. However, if the IP address points to a malicious server, this might cause the system to download malware.

Remember this

DNS poisoning attacks attempt to corrupt DNS data. A pharming attack redirects a web site's traffic to another web site and can do so by modifying the hosts file on the user's system.

ARP Poisoning Attacks

Address Resolution Protocol (ARP) poisoning is an attack that misleads computers or switches about the actual MAC address of a system. The MAC address is the physical address, or hardware address, assigned to the network interface card (NIC). ARP resolves the IP addresses of systems to their hardware address and stores the result in an area of memory known as the ARP cache.

TCP/IP uses the IP address to get a packet to a destination network. Once the packet arrives on the destination network, it uses the MAC address to get it to the correct host. ARP uses two primary messages:

- **ARP request.** The ARP request broadcasts the IP address and essentially asks, “Who has this IP address?”
- **ARP reply.** The computer with the IP address in the ARP request responds with its MAC address. The computer that sent the ARP request caches the MAC address for the IP. In many operating systems, all computers that hear the ARP reply also cache the MAC address.

A vulnerability with ARP is that it is very trusting. It will believe any ARP reply packet. Attackers can easily create ARP reply packets with spoofed or bogus MAC addresses, and poison the ARP cache on systems in the network. Two possible attacks from ARP poisoning are a man-in-the-middle attack and a DoS attack.

ARP Man-in-the-Middle Attacks

In a man-in-the-middle attack, an attacker can redirect network traffic, and in some cases insert malicious code. Consider Figure 7.3. Normally, traffic from the user to the Internet will go through the switch directly to the router, as shown in the top of Figure 7.3. However, after poisoning the ARP cache of the victim, traffic is redirected to the attacker.

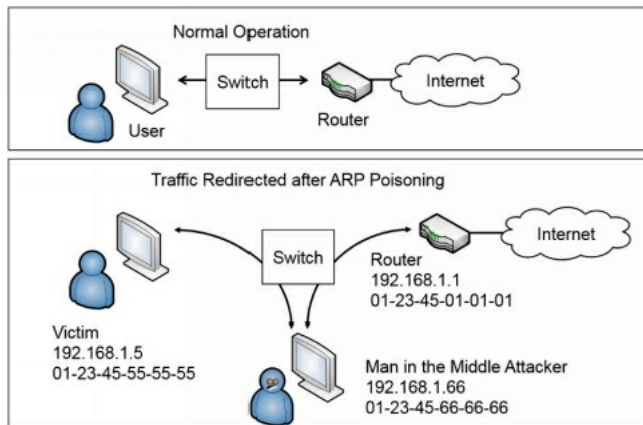


Figure 7.3: ARP poisoning used to redirect traffic

The victim’s ARP cache should include this entry to send data to the router:

192.168.1.1, 01-23-45-01-01-01

However, after poisoning the ARP cache, it includes this entry:

192.168.1.1, 01-23-45-66-66-66

The victim now sends all traffic destined for the router to the attacker. The attacker captures the data for analysis later. It also uses another method such as IP forwarding to send the traffic to the router so that the victim is unaware of the attack.

ARP DoS Attack

An attacker can also use ARP poisoning in a DoS attack. For example, an attacker can send an ARP reply with a bogus MAC address for the default gateway. The default gateway is the IP address of a router connection that provides a path out of the network. If all of the computers cache a bogus MAC address for the default gateway, none of them can reach it, and it stops all traffic out of the network.

Typo Squatting/URL Hijacking

Typo squatting (also called *URL hijacking*) occurs when someone buys a domain name that is close to a legitimate domain name. People often do so for malicious purposes. As an example, CompTIA hosts the comptia.org web site. If an attacker purchases the name comptai.org with a slight misspelling at the end of comptia, some users might inadvertently go to the attacker's web site instead of the legitimate web site.

Attackers might buy a similar domain for a variety of reasons, including:

- **Hosting a malicious web site.** The malicious web site might try to install drive-by malware on users' systems when they visit.
- **Earning ad revenue.** The attacker can host pay-per-click ads. When visitors click on the ads, advertisers pay revenue to the attacker.
- **Reselling the domain.** Attackers can buy domain names relatively cheaply, but resell them to the owner of the original site for a hefty profit.

Remember this

Attackers purchase similar domain names in typo squatting attacks for various malicious purposes. Users visit the typo squatting domain when they enter the URL incorrectly with a common typo.

Watering Hole Attacks

A *watering hole attack* attempts to discover which web sites employees are likely to visit and then infects those web sites with malware that can infect the visitors. The RSA Advanced Threat Intelligence Team first documented this attack in 2012.

Attackers apparently identified a number of web sites visited by personnel working in financial services. Next, they infected many of these web sites with malware that redirected the users to a malicious site. The malicious site attempted to install a type of remote access tool (RAT). When successful, the RAT allows attackers to remotely access and control infected systems.

Although one of the attacks focused on financial services, similar attacks have targeted other industries, including state and federal governments, educational institutions, and defense contractors. According to the RSA Advanced Threat Intelligence Team, approximately 32,000 users working in over 4,000 different organizations were redirected to the malicious web site.

Zero-Day Attacks

A *zero-day attack* is one that exploits an undocumented vulnerability. Many times, the vendor isn't aware of the issue. At some point, the vendor learns of the vulnerability and begins to write and test a patch to eliminate it. However, until the vendor releases the patch, the vulnerability is still a zero-day vulnerability.

As an example, a bug existed in the virtual DOS machine (VDM) that shipped with every version of 32-bit Windows systems from 1993 to 2010. The bug allowed attackers to escalate their privileges to full system level, effectively allowing them to take over the system. Google researcher Tavis Ormandy stated that he reported the bug to Microsoft in mid-2009. At this point, Microsoft (the vendor) knew about the bug, but didn't release a work-around until January 2010 and a patch until February 2010. Because the bug wasn't known publicly until January 2010, it remained a zero-day vulnerability until then.

Both attackers and security experts are constantly looking for zero-day vulnerabilities. Attackers want to learn about them so that they can exploit them. Most security experts want to know about them so that they can help ensure that vendors patch them before causing damage to users.

Remember this

Zero-day exploits are undocumented and unknown to the public. The vendor might know about it, but has not yet released a patch to address it.

Web Browser Concerns

Users surf the Internet with web browsers. In the context of security, there are some issues related to web browsers that cause some problems and this section addresses many common concerns.

Malicious Add-Ons

Many web browsers support add-ons to enhance the capability of the browser. For example, you can install the Adobe PDF reader add-on into a browser to automatically open PDF files within the browser window. Similarly, some add-ons include pop-up blockers to prevent these pop-ups from appearing.

Although many add-ons are helpful, some are malicious. As an example, the Mozilla Sniffer add-on added malicious capabilities to the Firefox browser. After installation, it intercepted the user's logon data submitted to any web site and sent it to a remote location, presumably managed by attackers. The add-on was only available for a short time as an experimental add-on, but was downloaded and installed by at least 1,800 users. Users should be cautious when installing new add-ons.

Cookies and Attachments

A *cookie* is a text file stored on a user's computer and used for multiple purposes, including tracking a user's activity. Web sites regularly write cookies on user systems to help remember the user and enhance the user experience.

As an example, Amazon makes frequent use of cookies. When I visit the site and look at different products, it tracks my activity and places ads on the web site based on my previous searches or purchases. In most cases, only the web site can read the cookie. However, cross-site scripting attacks (described later in this chapter) allow attackers to read cookies.

Some web developers store sensitive data, such as usernames or passwords, in cookies. If attackers can read the cookies, they may have access to sensitive data. Additionally, cookies include a session ID that can identify the user session when the user logs on, and this session ID can be used in a session hijacking attack.

Attachments are typically associated with emails. For example, if you want to share a file with someone else, you can attach the file to your email and send it. Attackers often use attachments when sending malicious spam. If the user opens the attachment, it attempts to install malware onto the user's system. As an example, I received a malicious email today with a PDF file. The subject was "Notice of court attendance." It indicated I was scheduled to attend a court hearing and encouraged me to thoroughly study the plaintiff note in the attachment. The attachment is a Zip file that includes malware, which installs itself just by opening the Zip file.

Session Hijacking Attacks

When a user logs on to a web site, the web site often returns a cookie with a session ID. In many cases, this cookie is stored on the user's system and remains active until the user logs off. If the user closes the session and returns to the web site, the web site reads the cookie and automatically logs the user on. This is convenient for the user, but can be exploited by an attacker.

In a *session hijacking* attack, the attacker learns the user's session ID and uses it to impersonate the user. The web server doesn't know the difference between the original user and the attacker because it is only identifying the user based on the session ID.

Attackers can read cookies installed on systems through several methods, including cross-site scripting attacks and Flash cookies (described in the next section). Once they have the session ID, they can use header manipulation to hijack the session.

Flash Cookies and LSOs

A *Flash cookie* is one created by Adobe Flash Player and is different from a traditional text cookie. They are also known as *local shared objects* (*LSOs*) or locally shared objects. As one example, Flash cookies are stored in multiple locations by default, and traditional methods of deleting cookies through a web browser do not delete Flash cookies. Some Flash cookies store the session ID from traditional cookies. If a user deletes the traditional cookies, the Flash cookies re-create them.

Many sites use Flash cookies to track users' online activity without their knowledge or consent. For example, when a user goes to a web site using a Flash cookie, then goes to web site B, the Flash cookie records their activity on web site B. This continues for the entire session as the Flash cookie tracks and records all of the user's activity. Their usage has prompted many class-action lawsuits against sites using Flash cookies.

Arbitrary Code Execution/Remote Code Execution

Arbitrary code execution refers to the ability of an attacker to execute commands or run programs on a target system. *Remote code execution* refers to the ability of an attacker to execute the code from a remote system. Neither of these is desirable because it allows attackers to install and run malware

on vulnerable systems.

As an example, imagine an application such as a web browser has a vulnerability that allows execution of code using one of these methods. An unsuspecting user can visit a malicious web site with specially crafted code to exploit this vulnerability. This code can use elevated privileges to cause the system to download a malicious file and then execute it to install it. Once installed, the malware can allow the attacker to take control of the computer whenever desired.

Software bugs are the most common reason that arbitrary code execution and remote code execution is possible. This is another reason why it is important to keep systems up to date with current patches.

Header Manipulation Attacks

TCP/IP packages data into packets before sending them over a network. These packets have headers, which include different types of information depending on the header type. For example, TCP headers include port numbers to identify the protocol, and IP headers include source and destination IP addresses.

Headers also include various flags. A *flag* is simply a bit that is set to a 1 or a 0, often indicating true or false. Attackers can manipulate the flags within the headers to modify behavior. In some cases of header manipulation, the attacker modifies data within the packet, such as the session ID. Many programs are available to attackers, making it relatively easy to modify these headers.

In a session hijacking attack, the attacker inserts the session ID of the original user into the header. If the web server uses this session ID to log the user on automatically, it gives the attacker access to the user's account.

Many web sites use dual authentication to prevent an attacker from taking malicious action with the session ID. For example, Amazon will use the session ID to identify the user and enhance the browsing experience. However, if the user makes a purchase, Amazon requires the user to authenticate again.

Understanding Secure Coding Concepts

Applications often provide a method for attackers to generate attacks unless developers create them using secure coding concepts. These include input validation and error handling, both covered in this section.

Although it's important to ensure that all applications are secure, the CompTIA Security+ exam focuses on server applications, such as web server applications hosted on the Internet. Chapter 5, "Securing Hosts and Data," presents many of the concepts used to harden operating systems and applications. As a reminder, some common hardening steps include disabling unnecessary services and keeping systems up to date. This section includes some additional steps.

Performing Input Validation

One of the most important security steps that developers should take is to include input validation. Input validation is the practice of checking data for validity before using it. Input validation prevents an attacker from sending malicious code that an application will use by either sanitizing the input to remove malicious code or rejecting the input. The lack of input validation is one of the most common security issues on web-based applications. It allows many different types of attacks, such as buffer overflow, SQL injection, command injection, and cross-site scripting attacks (covered in the next section).

Consider a web form that includes a text box for a first name. You can logically expect a valid first name to have only letters, and no more than 25 letters. The developer uses input validation techniques to ensure that the name entered by the user meets this validity check. If a user enters other data, such as numbers, semicolons, or Hypertext Markup Language (HTML) code, it fails the validity check. Instead of using the data, the application rejects it and provides an error to the user.

You've probably seen input validation checks and error-handling routines in use if you've ever filled out a form on a web page. If you didn't fill out all the required text boxes, or if you entered invalid data into one or more of the boxes, the web site didn't crash. Instead, it redisplayed the page and showed an error. Web sites often use a red asterisk next to text boxes with missing or invalid data.

Some common checks performed by input validation include:

- **Verifying proper characters.** Some fields such as a zip code use only numbers, whereas other fields such as state name use only letters. Other fields are a hybrid. For example, a phone number uses only numbers and dashes. Developers can configure input validation code to check for specific character types, and even verify that characters are entered in the correct order. For example, a telephone number mask of ###-###-#### accepts only three numbers, a dash, three numbers, a dash, and four numbers.
- **Implementing boundary or range checking.** These checks ensure that values are within expected boundaries or ranges. For example, if the maximum purchase for a product is three, a range check verifies the quantity is three or less. The validation check identifies data outside the range as invalid and the application does not use it.
- **Blocking HTML code.** Some malicious attacks embed HTML code within the input as part of an attack. These can be blocked by preventing the user from entering the < and > characters, which are used within HTML code.
- **Preventing the use of certain characters.** Some attacks, such as SQL injection attacks, use specific characters such as the dash (-), apostrophe ('), and equal sign (=). Blocking these characters helps to prevent these attacks.

Client-Side and Server-Side Input Validation

It's possible to perform input validation at the client and the server. Client-side input validation is quicker, but is vulnerable to attacks. Server-side input validation takes longer, but is secure because it ensures the application doesn't receive invalid data. Many applications use both. Imagine Homer is using a web browser to purchase the newest version of Scrabble ships through the Duff web site. Customers cannot purchase more than three at a time.

In client-side input validation, the validation code is included in the HTML page sent to Homer. If he enters a quantity of four or more, the HTML code gives him an error message, and doesn't submit the page to the server until Homer enters the correct data.

Unfortunately, it's possible to bypass client-side validation techniques. Many web browsers allow users to disable JavaScript in the web browser, which bypasses client-side validation. It's also possible to use a web proxy to capture the data sent from the client in the HTTP POST command and modify it before forwarding to the server.

Server-side input validation checks the inputted values when it reaches the server. This ensures that the user hasn't bypassed the client-side checks.

Using both client-side and server-side validation provides speed and security. The client-side validation checks prevent round-trips to the user until the user has entered the correct data.

Remember this

The lack of input validation is one of the most common security issues on web-based applications. Input validation verifies the validity of inputted data before using it, and server-side validation is more secure than client-side validation. Input validation protects against many attacks, such as buffer overflow, SQL injection, command injection, and cross-site scripting attacks.

Avoiding Race Conditions

When two or more modules of an application, or two or more applications, attempt to access a resource at the same time, it can cause a conflict known as a race condition. Most application developers are aware of race conditions and include methods to avoid them when writing code. However, when new developers aren't aware of race conditions, or they ignore them, a race condition can cause significant problems.

As a simple example of a potential problem, imagine you are buying a plane ticket online and use the web application to pick your seat. You find a window seat and select it. However, at the same time you're selecting this window seat, someone else is, too. You both make the purchase at the same

time and you both have tickets with the same seat number. You arrive after the other person and he's unwilling to move, showing his ticket with the seat number. A flight attendant ultimately helps you find a seat. Unfortunately, it's between two burly gentlemen who have been on an all-cabbage diet for the last week. You probably wouldn't be too happy.

In reality, online ticketing applications for planes, concerts, and more avoid this type of race condition. In some cases, they lock the selection before offering it to a customer. In other cases, they double-check for a conflict later in the process. Most database applications have internal concurrency control processes to prevent two entities from modifying a value at the same time. However, inexperienced web application developers often overlook race conditions.

Error and Exception Handling

Error and exception handling routines are a part of input validation, and they ensure that an application can handle an error gracefully. They catch errors and provide user-friendly feedback to the user. When an application doesn't catch an error, it can cause the application to fail. In the worst-case scenario, a faulty application can cause the operating system to crash. Using effective error- and exception-handling routines protects the integrity of the underlying operating system.

When an application doesn't catch an error, it often provides debugging information that attackers can use against the application. In contrast, when an application catches the error, it can control what information it shows to the user. There are two important points about error reporting:

- **Errors to users should be general.** Detailed errors provide information that attackers can use against the system, so the errors should be general. Attackers can analyze the errors to determine details about the system. For example, if an application is unable to connect with a database, the returned error can let the attacker know exactly what type of database the system is running. This indirectly lets the attacker know what types of commands the system will accept. Also, detailed errors confuse most users.
- **Detailed information should be logged.** Detailed information on the errors typically includes debugging information. This information makes it easier for developers to identify what caused the error and how to resolve it.

Remember this

Error and exception handling helps protect the integrity of the operating system and controls the errors shown to users. Applications should show generic error messages to users but log detailed information.

Identifying Application Attacks

Many attacks target server applications such as those hosted on web servers. Web servers are highly susceptible to several types of attacks, such as buffer overflow attacks and SQL injection attacks, because they commonly accept data from users. Other servers are susceptible to some types of command injection attacks. This section covers many of the common attacks related to different types of servers.

Web Servers

Web servers most commonly host web sites accessible on the Internet, but they can also serve pages within an internal network. Organizations place web servers within a demilitarized zone (DMZ) to provide a layer of protection.

The two primary applications used for web servers are:

- **Apache.** Apache is the most popular web server used on the Internet. It's free and can run on Unix, Linux, and Windows systems.
- **Internet Information Services (IIS).** IIS is a Microsoft web server, and it's included free with any Windows Server product.

Establishing a web presence is almost a requirement for organizations today, and users expect fancy web sites with dynamic pages that are easy to use. Although many applications make it easy to create web sites, they don't always include security. This often results in many web sites being highly susceptible to attacks. The following sections identify many common attacks on web sites.

Buffer Overflows and Buffer Overflow Attacks

A *buffer overflow* occurs when an application receives more input, or different input, than it expects. The result is an error that exposes system memory that would otherwise be protected and inaccessible. Normally, an application will have access only to a specific area of memory, called a buffer. The buffer overflow allows access to memory locations beyond the application's buffer, enabling an attacker to write malicious code into this area of memory.

As an example, an application may be expecting to receive a string of 15 characters for a username. If input validation is not used and it receives more than 15 characters, it can cause a buffer overflow and expose system memory. The following HTTP GET command shows an example of sending a long string to the system to create a buffer overflow:

```
GET /index.php?username=////////////////////////////////////
```

The buffer overflow exposes a vulnerability, but it doesn't necessarily cause damage by itself. However, once attackers discover the vulnerability, they exploit it and overwrite memory locations with their own code. If the attacker uses the buffer overflow to crash the system or disrupt its services, it is a DoS attack.

More often, the attacker's goal is to insert malicious code in a memory location that the system will execute. It's not easy for an attacker to know the exact memory location where the malicious code is stored, making it difficult to get the computer to execute it. However, an attacker can make educated guesses to get close.

A popular method that makes guessing easier is with no operation (NOP, pronounced as "no-op") commands, written as a NOP slide or NOP sled. Many Intel processors use hexadecimal 90 (often written as x90) as a NOP command, so a string of x90 characters is a NOP sled. The attacker writes a long string of x90 instructions into memory, followed by malicious code. When a computer is executing code from memory and it comes to a NOP, it just goes to the next memory location. With a long string of NOPs, the computer simply slides through all of them until it gets to the last one and then executes the code in the next instruction. If the attacker can get the computer to execute code from a memory location anywhere in the NOP slide, the system will execute the attacker's malicious code.

The malicious code varies. In some instances, the attackers write code to spread a worm through the web server's network. In other cases, the code modifies the web application so that the web application tries to infect every user who visits the web site with other malware. The attack possibilities are almost endless.

Remember this

Buffer overflows occur when an application receives more data than it can handle, or receives unexpected data that exposes system memory. Buffer overflow attacks often include NOP instructions (such as x90) followed by malicious code. When successful, the attack causes the system to execute the malicious code. Input validation helps prevent buffer overflow attacks.

A buffer overflow attack includes several different elements, but they happen all at once. The attacker sends a single string of data to the application. The first part of the string causes the buffer overflow. The next part of the string is a long string of NOPs followed by the attacker's malicious code, stored in the attacked system's memory. Last, the malicious code goes to work.

In some cases, an attacker is able to write a malicious script to discover buffer overflow vulnerabilities. For example, the attacker could use JavaScript to send random data to another service on the same system.

Although error-handling routines and input validation go a long way to prevent buffer overflows, they don't prevent them all. Attackers occasionally discover a bug allowing them to send a specific string of data to an application causing a buffer overflow. When vendors discover buffer overflow vulnerabilities, they are usually quick to release a patch or hotfix. From an administrator's perspective, the solution is easy: Keep the systems up to date with current patches.

Integer Overflow

An *integer overflow attack* attempts to create a numeric value that is too big for an application to handle. The result is that the application gives inaccurate results. As an example, if an application reserves eight bits to store a number, it can store any value between 0 and 255. If the application attempts to multiply two values such as 95×59 , the result is 5,605. This number cannot be stored in the eight bits, so it causes an integer overflow error.

In some situations, an integer overflow error occurs if an application expects a positive number, but receives a negative number instead. If the application doesn't have adequate error- and exception-handling routines, this might cause a buffer overflow error.

SQL Queries and SQL Injection Attacks

SQL (pronounced as “sequel” or “es-que-el”) is a Structured Query Language used to communicate with databases. SQL statements read, insert, update, and delete data to and from a database. Many web sites use SQL statements to interact with a database providing users with dynamic content.

The following sections identify how SQL queries work, how attackers launch a SQL injection attack, and how to protect against SQL injection attacks.

SQL Queries

As a simple example of a web site that uses SQL queries, think of Amazon.com. When you enter a search term and click Go (as shown in Figure 7.4), the web application creates a SQL query, sends it to a database server, and formats the results into a web page that it sends back to you.



Figure 7.4: Web page querying a database with SQL

In the example, I selected the Books category and entered **Darril Gibson**. The result shows a list of books authored by Darril Gibson available for sale on Amazon. The query sent to the database from the Amazon web application may look like this:

```
SELECT * FROM Books WHERE Author = 'Darril Gibson'
```

The * is a wildcard and returns all columns in a table. Notice that the query includes the search term entered into the web page form (Darril Gibson) and encloses the search term in single quotes. If the web site simply plugs the search term into the SELECT statement, surrounded by single quotes, it will work, but it’s also highly susceptible to SQL injection attacks.

SQL Injection Attacks

In a SQL injection attack, the attacker enters additional data into the web page form to generate different SQL statements. SQL query languages use a semicolon (;) to indicate the end of the SQL line and use two dashes (--) as an ignored comment. With this knowledge, the attacker could enter different information into the web form like this:

```
Darril Gibson'; SELECT * FROM Customers;--
```

If the web application plugged this string of data directly into the SELECT statement surrounded by the same single quotes, it would look like this:

```
SELECT * FROM Books WHERE Author = 'Darril Gibson';  
SELECT * FROM Customers;  
--'
```

The first line retrieves data from the database, just as before. However, the semicolon signals the end of the line and the database will accept another command. The next line reads all the data in the Customers table, which can give the attacker access to names, credit card data, and more. The last line comments out the second single quote to prevent a SQL error.

If the application doesn’t include error-handling routines, these errors provide details about the type of database the application is using, such as an Oracle, Microsoft SQL Server, or MySQL database. Different databases format SQL statements slightly differently, but once the attacker learns the database brand, it’s a simple matter to format the SQL statements required by that brand. The attacker then follows with SQL statements to access the database and may allow the attacker to read, modify, delete, and/or corrupt data.

This attack won’t work against Amazon (please don’t try it) because Amazon is using secure coding principles. I don’t have access to its code, but I’d bet their developers are using input validation and SQL-based stored procedures.

Many SQL injection attacks use a phrase of **or '1' = '1'** to create a true condition. For example, if an online database allows you to search a Customers table looking for a specific record, it might expect you to enter a name. If you entered **Homer Simpson**, it would create a query like this:

```
SELECT * FROM Customers WHERE name = 'Homer Simpson'
```

This query will retrieve a single record for Homer Simpson. However, if the attacker enters **' or '1'='1' --** instead of Homer Simpson, it will create a query like this:

```
SELECT * FROM Customers WHERE name = '' or '1'='1' --'
```

Although this is a single SELECT statement, the **or** clause causes it to behave as two separate SELECT statements:

```
SELECT * FROM Customers WHERE name = ''  
SELECT * FROM Customers WHERE '1'='1'
```

The first clause will likely not return any records because the table is unlikely to have any records with the name field empty. However, because the number 1 always equals the number 1, the WHERE clause in the second statement always equates to True, so the SELECT statement retrieves all records from the Customers table.

In many cases, a SQL injection attack starts by sending improperly formatted SQL statements to the system to generate errors. Proper error handling

prevents the attacker from gaining information from these errors, though. Instead of showing the errors to the user, many web sites simply present a generic error web page that doesn't provide any details.

Remember this

Attackers use SQL injection attacks to pass queries to back-end databases through web servers. Many SQL injection attacks use the phrase ' or '1'='1' – to trick the database server into providing information. Input validation and stored procedures reduce the risk of SQL injection attacks.

Protecting Against SQL Injection Attacks

As mentioned previously, input validation provides strong protection against SQL injection attacks. Before using the data entered into a web form, the web application verifies that the data is valid.

Additionally, database developers often use stored procedures with dynamic web pages. A stored procedure is a group of SQL statements that execute as a whole, similar to a mini-program. A parameterized stored procedure accepts data as an input called a parameter. Instead of copying the user's input directly into a SELECT statement, the input is passed to the stored procedure as a parameter. The stored procedure performs data validation, but it also handles the parameter (the inputted data) differently and prevents a SQL injection attack.

Consider the previous example searching for a book by an author where an attacker entered the following text: **Darril Gibson'; SELECT * From Customers;--**. The web application passes this search string to a stored procedure. The stored procedure then uses the entire search string in a SELECT statement like this:

```
SELECT * From Books Where Author = "Darril Gibson'; SELECT * From Customers;-- "
```

In this case, the text entered by the user is interpreted as harmless text rather than malicious SQL statements. It will look for books with an author name using all of this text: Darril Gibson'; SELECT * From Customers;--. People don't have names with SELECT statements embedded in them so the query comes back empty.

Depending on how well the database server is locked down (or not), SQL injection attacks may allow the attacker to access the structure of the database, all the data, and even modify data. In some cases, attackers have modified the price of products from several hundred dollars to just a few dollars, purchased several of them, and then returned the price to normal.

XML Injection

Many databases use Extensible Markup Language (XML) for inputting or exporting data. XML provides formatting rules to describe the data. For example, here's an XML tag for a name: <name>Darril Gibson</name>. The data is "Darril Gibson" and the XML tags (<name> and </name>) describe the data as a name.

Additionally, databases use XPath as a query language for XML data. If an application accepts XML data without input validation and without stored procedures, it is susceptible to an XML injection attack similar to a SQL injection attack. The attacker can insert additional data in an XML injection attack. This additional data creates XPath statements to retrieve or modify data.

NoSQL Versus SQL Databases

Server-based SQL databases are traditional relational databases using tables that relate to each other in one way or another. They are very effective in many situations, but not all. A newer type of database has emerged known as not only SQL (NoSQL).

NoSQL databases typically hold one or more of the following types of data: documents, key-value pairs, or graphs. Documents are formatted in a specific way and each document represents an object. This is similar to how a table holds data in rows. However, the document-based NoSQL database gives developers much more flexibility in how they can store and query the data.

Both NoSQL and SQL databases are susceptible to command injection attacks if developers do not implement input validation techniques. SQL databases use SQL queries and are susceptible to SQL injection attacks. NoSQL databases use unstructured query language (UQL) queries. Although the format of UQL queries varies with different vendors, attackers can learn them and use them when developers do not implement input validation techniques

Cross-Site Scripting

Cross-site scripting (XSS) is another web application vulnerability that can be prevented with input validation. Attackers embed malicious HTML or JavaScript code into an email or web site error message. If a user responds to the email or error message, it executes the code. Many times, this gives the attacker access to user cookies or other information about the user.

You may be wondering why the acronym isn't CSS instead of XSS. The reason is the web sites use Cascading Style Sheets identified as CSS and CSS files are not malicious.

HTML and JavaScript code use tags surrounded by the less-than (<) and greater-than (>) characters. For example, images are placed within a web page with the tag. These tags support many additional options and commands, which developers use to create feature-rich web pages. However, attackers can manipulate these tags to run malicious code. Attackers often embed cross-scripting code into comments on blog pages or forums when the page allows users to include HTML tags.

As an example, a bug in Twitter's web site resulted in a cross-site scripting problem in 2010. A malicious Twitter user discovered the bug and embedded JavaScript code into a tweet. When innocent users opened a web page that included the malicious tweet, the code ran on their systems. In this case, the malicious code used the onMouseOver event. When users hovered their mouse over the tweet, it did two things. It retweeted the tweet, sending it out to all of the user's followers. It also launched a pop-up window displaying content from a hard-core Japanese pornography web site.

Although this attack was more embarrassing to Twitter than harmful to end users, many other cross-site scripting attacks are malicious. They can allow attackers to redirect users to other web sites, steal cookies off a user's system, read passwords from a web browser's cache, and more. If a web site stored private data in a user's cookie, such as a username and password, an attacker can use a cross-site scripting attack to retrieve this information.

The primary protection against cross-site scripting attacks is at the web application with input validation techniques to block the use of HTML tags and JavaScript tags. Tags are embedded within the < and > characters, so it's possible to block these tags by rejecting any text that includes these characters. It's also important to educate users about the dangers of clicking links. Some XSS attacks send emails with malicious links within them. The XSS attack fails if users do not click the link.

Remember this

Cross-site scripting (XSS) attacks allow attackers to capture user information such as cookies. Input validation techniques at the server help prevent XSS attacks.

Cross-Site Request Forgery (XSRF)

Cross-site request forgery (XSRF or CSRF) is an attack where an attacker tricks a user into performing an action on a web site. The attacker creates a specially crafted HTML link and the user performs the action without realizing it.

As an innocent example of how HTML links create action, consider this HTML link: <http://www.google.com/search?q=Success>. If a user clicks on the link, it works just as if the user browsed to Google and entered Success as a search term. The *?q=Success* part of the query causes the action.

Many web sites use the same type of HTML queries to perform actions. For example, imagine a web site that supports user profiles. If users wanted to change profile information, they could log on to the site, make the change, and click a button. The web site may use a link like this to perform the action:

<http://getcertifiedgetahead.com/edit?action=set&key=email&value=you@home.com>

Attackers use this knowledge to create a malicious link. For example, the following link could change the email address in the user profile, redirecting the user's email to the attacker:

<http://getcertifiedgetahead.com/edit?action=set&key=email&value=hacker@hackersrs.com>

Although this shows one possibility, there are many more. If a web site supports any action via an HTML link, an attack is possible. This includes making purchases, changing passwords, transferring money, and much more.

Web sites typically won't allow these actions without users first logging on. However, if users have logged on before, authentication information is stored on their system either in a cookie or in the web browser's cache. Some web sites automatically use this information to log users on as soon as they visit. In some cases, the XSRF attack allows the attacker to access the user's password.

Users should be educated on the risks related to links from sources they don't recognize. Phishing emails (covered in Chapter 6, "Understanding Malware and Social Engineering") often include malicious links that look innocent enough to users, but can cause significant harm. If users don't click the link, they don't launch the XSRF attack.

However, just as with cross-site scripting, the primary burden of protection from XSRF falls on the web site developers. Developers need to be aware of XSRF attacks and the different methods used to protect against them. One method is to use dual authentication and force the user to manually enter credentials prior to performing actions. Another method is to expire the cookie after a short period, such as after 10 minutes, preventing automatic logon for the user.

Remember this

Cross-site request forgery (XSRF) scripting causes users to perform actions on web sites, such as making purchases, without their knowledge. In some cases, it allows an attacker to steal cookies and harvest passwords.

Directory Traversal/Command Injection

In some cases, attackers are able to inject operating system commands into an application using web page forms or text boxes. Any web page that accepts input from users is a potential threat. Directory traversal is a specific type of command injection attack that attempts to access a file by including the full directory path, or traversing the directory structure.

For example, in Unix systems, the `passwd` file includes user logon information, and it is stored in the `/etc` directory with a full directory path of `/etc/passwd`. Attackers can use commands such as `../../etc/passwd` or `/etc/passwd` to read the file. Similarly, they could use a remove directory command (such as `rm -rf`) to delete a directory, including all files and subdirectories. Input validation can prevent these types of attacks.

LDAP Injection

Chapter 3 introduced Lightweight Directory Application Protocol (LDAP). As a reminder, this is the primary protocol used to communicate with servers hosting directory services such as Active Directory in a Microsoft system. In some cases, attackers are able to use LDAP injection attacks to query and modify account information in Active Directory with LDAP commands.

Remember this

LDAP injection attacks attempt to access or modify data hosted on directory service servers.

Transitive Access and Client-Side Attacks

Transitive access relationships refer to trusts, and if not addressed, these relationships can allow unauthorized personnel to access restricted data. SQL injection attacks are a type of transitive access attack performed from the client side using a web browser.

As an example, consider Figure 7.5. Homer is able to access the web server, and the web server is able to access the database server. However, even though Homer is not able to access the database server directly, he might be able to access the database server using transitive access via his web browser.




Figure 7.5: Protected database server

The following security controls are in place:

- **FW1.** Rules in the ACL of FW1 open ports 80 and 443 allowing HTTP and HTTP Secure (HTTPS traffic), respectively.
- **FW2.** Microsoft SQL Server uses port 1433, so a rule in the ACL of FW2 allows traffic from the web server to the database server over port 1433. FW2 blocks all other traffic using port 1433.
- **Encryption of data at rest.** Encryption on the database server protects confidential customer data such as credit card information. Administrators commonly encrypt the fields holding this data, but don't encrypt the entire database. This protects the data if thieves steal the database server, or if an attacker copies the entire database file.
- **Trust relationship between web server and database server.** This trust relationship allows the web server to query the database server and retrieve data. This trust relationship also ensures the database application decrypts data at rest before sending it to the web server.
- **Data in transit between the web server and the database server.** Transport Layer Security (TLS) encrypts data in transit between the web server and database server. This prevents an attacker from reading data captured with a protocol analyzer in a sniffing attack.
- **Data in transit between the customer and the web server.** HTTPS encrypts sensitive data between the customer and the web server. This also prevents sniffing attacks.

All of these security controls provide a strong defense-in-depth strategy with several layers of security. However, they aren't enough.

If the application on the web server isn't using input validation techniques, an attacker can use a SQL injection attack from the client side. The SQL injection attack uses transitive access to get to the database server through the web server.

In addition to explaining a client-side attack using transitive access, this also provides a great reminder that security is never done. Network and database administrators have implemented a strong defense-in-depth strategy. However, if web application developers don't implement security controls such as input validation, attackers can bypass all of the security controls and access the data.

Remember this

A client-side attack uses an application on the client computer, such as a web browser. A transitive access attack attempts to access a back-end server through another server. A SQL injection attack is an example of a transitive access attack that can bypass many other security controls.

Fuzzing

Fuzzing (or fuzz testing) uses a computer program to send random data to an application. In some cases, the random data can actually crash the program or provide unexpected results, indicating a vulnerability. Security professionals use fuzz testing to test systems and applications for vulnerabilities they can correct. In contrast, attackers use fuzz testing to identify vulnerabilities they can exploit.

Sometimes attackers will write a fuzz testing script to run on the attacked system instead of sending the data over the network. For example, an attacker can use JavaScript to send random data to another service on the same system. In some cases, this discovers a string of code that can cause a buffer overflow. If an attacker discovers a string of data that can create a buffer overflow, he can use it in an attack.

Remember this

Fuzzing sends random strings of data to applications looking for vulnerabilities. Administrators use fuzz testing to test applications and attackers use fuzzing to detect attack methods.

Chapter 7 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Comparing Common Attacks

- A DoS attack is an attack launched from a single system and attempts to disrupt services.
- DDoS attacks are DoS attacks from multiple computers. DDoS attacks typically include sustained, abnormally high network traffic.
- Smurf attacks spoof the source IP address and use a directed broadcast ping to flood victims with ping replies. Smurf attacks often use amplifying networks. Configuring routers to block directed broadcasts prevents a network from becoming an amplifying network.
- Replay attacks capture data in a session with the intent of using information to impersonate one of the parties. Timestamps and sequence numbers thwart replay attacks.
- Account lockout policies thwart online password attacks such as dictionary and brute force attacks that attempt to guess a password. Complex passwords thwart offline password attacks.
- Password salting adds additional characters to passwords before hashing them, and prevents many types of attacks, including dictionary, brute force, and rainbow table attacks.
- DNS poisoning attacks modify DNS data and can redirect users to malicious sites. Pharming attacks often modify the hosts file to redirect web site traffic to a malicious web site.
- Attackers buy domain names with minor typographical errors in the hopes of attracting traffic when users enter the incorrect URL. Attackers can configure the sites with malware to infect visitors or configure the site to generate ad revenue for the attacker.
- Attackers exploiting unknown or undocumented vulnerabilities are taking advantage of zero-day vulnerabilities. The vulnerability is no longer a zero-day vulnerability after the vendor releases a patch to fix it.

Understanding Secure Coding Concepts

- A common coding error in web-based applications is the lack of input validation.
- Input validation checks the data before passing it to the application and prevents many types of attacks, including buffer overflow, SQL injection, command injection, and cross-site scripting attacks.
- Server-side input validation is the most secure. Attackers can bypass client-side input validation, but not server-side input validation.
- Error- and exception-handling routines within applications can prevent application failures and protect the integrity of the operating systems. Error messages shown to users should be generic, but the application should log detailed information on the error.

Identifying Application Attacks

- Buffer overflows occur when an application receives more data, or unexpected data, than it can handle and exposes access to system memory.
- Buffer overflow attacks exploit buffer overflow vulnerabilities. A common method uses NOP instructions or NOP sleds such as a string of x90 commands. Two primary protection methods against buffer overflow attacks are input validation and keeping a system up to date.
- SQL injection attacks provide information about a database and can allow an attacker to read and modify data within a database from a web page. Input validation and stored procedures provide the best protection.
- Client-side attacks originate from the client such as within a web browser. Transitive access attacks attempt to access resources via a transitive trust relationship. SQL injection is an example of a client-side transitive access attack.
- Cross-site scripting (XSS) allows an attacker to redirect users to malicious web sites and steal cookies. It uses HTML and JavaScript tags with < and > characters.
- Cross-site request forgery (XSRF) causes users to perform actions on web sites without their knowledge and allows attackers to steal cookies and harvest passwords.
- XSS and XSRF attacks are mitigated with input validation techniques.
- Lightweight Directory Application Protocol (LDAP) injection attacks attempt to access data on servers hosting a directory service, such as Microsoft Active Directory.
- Transitive access attacks can attack back-end servers via a front-end server. For example, SQL injection attacks start as a client-side attack, but access back-end databases via a web server.
- Fuzzing sends random data to an application to test the application's ability to handle the random data. Fuzzing can cause an application to crash if proper input validation techniques are not used.

Chapter 7 Practice Questions

1. An IDS alerts on increased traffic. Upon investigation, you realize it is due to a spike in network traffic from several sources. Assuming this is malicious, what is the MOST likely explanation?
 - A. A smurf attack
 - B. A flood guard attack
 - C. A DoS attack
 - D. A DDoS attack
2. A network administrator needs to ensure the company's network is protected against smurf attacks. What should the network administrator do?
 - A. Install flood guards.
 - B. Use salting techniques.
 - C. Verify border routers block directed broadcasts.
 - D. Ensure protocols use timestamps and sequence numbers.
3. Some protocols include timestamps and sequence numbers. These components help protect against what type of attacks?
 - A. Smurf
 - B. Replay
 - C. Flood guards
 - D. Salting
4. Which of the following is the BEST method to protect against someone trying to guess the correct PIN to withdraw money from an ATM?
 - A. Account lockout
 - B. Rainbow table
 - C. Salting
 - D. Input validation
5. An application stores user passwords in a hashed format. Which of the following can decrease the likelihood that attackers can discover these passwords?
 - A. Rainbow tables
 - B. MD5
 - C. Salt
 - D. Smurf
6. A user complains that his system is no longer able to access the `blogs.getcertifiedgetahead.com` site. Instead, his browser goes to a different site. After investigation, you notice the following entries in the user's hosts file:
127.0.0.1 localhost
72.52.230.233 `blogs.getcertifiedgetahead.com`
What is the BEST explanation for this entry?
 - A. A pharming attack
 - B. A whaling attack
 - C. Session hijacking
 - D. A phishing attack
7. Security analysts recently discovered that users in your organization are inadvertently installing malware on their systems after visiting the `comptai.org` web site. Users have a legitimate requirement to visit the `comptia.org` web site. What is the MOST likely explanation for this activity?
 - A. Smurf
 - B. Typo squatting
 - C. Fuzzing
 - D. Replay

8. An attacker recently attacked a web server hosted by your company. After investigation, security professionals determined that the attacker used a previously unknown application exploit. Which of the following BEST identifies this attack?
- A. Buffer overflow
 - B. Zero-day attack
 - C. Fuzzing
 - D. Session hijacking
9. Which of the following developer techniques results in significant security vulnerabilities for online web site applications?
- A. Buffer overflow
 - B. XSRF
 - C. Poor input validation
 - D. Hardening
10. An attacker is bypassing client-side input validation by intercepting and modifying data within the HTTP POST command. Which of the following does the attacker use in this attack?
- A. Command injection
 - B. Flash cookie
 - C. Proxy
 - D. Exception handling
11. Web developers are implementing error and exception handling in a web site application. Which of the following represents a best practice for this?
- A. Displaying a detailed error message but logging generic information on the error
 - B. Displaying a generic error message but logging detailed information on the error
 - C. Displaying a generic error message and logging generic information on the error
 - D. Displaying a detailed error message and logging detailed information on the error
12. While reviewing logs for a web application, a developer notices that it has crashed several times reporting a memory error. Shortly after it crashes, the logs show malicious code that isn't part of a known application. What is MOST likely occurring?
- A. Buffer overflow
 - B. XSS
 - C. Cross-site scripting
 - D. XML injection
13. An application on one of your database servers has crashed several times recently. Examining detailed debugging logs, you discover that just prior to crashing, the database application is receiving a long series of x90 characters. What is MOST likely occurring?
- A. SQL injection
 - B. Buffer overflow
 - C. XML injection
 - D. Zero-day
14. Attackers have attacked an online web server using a SQL injection attack. Which of the following BEST describes this?
- A. The attacker is attempting to overload the system with unexpected data and access memory locations.
 - B. The attacker is attempting to impersonate a user using HTML code.
 - C. The attacker is sending random data into a program to see if the application will crash.
 - D. The attacker is attempting to pass commands to a back-end database server to access data.
15. While creating a web application, a developer adds code to limit data provided by users. The code prevents users from entering special characters. Which of the following attacks will this code MOST likely prevent?
- A. Sniffing
 - B. Spoofing
 - C. XSS
 - D. Pharming

16. Homer recently received an email thanking him for a purchase that he did not make. He asked an administrator about it and the administrator noticed a pop-up window, which included the following code:

```
<body onload="document.getElementById('myform').submit()">
  <form id="myForm" action="gcgapremium.com/purchase.php" method="post"
    <input name="Buy Now" value="Buy Now" />
  </form>
</body>
```

What is the MOST likely explanation?

- A. XSRF
 - B. Buffer overflow
 - C. SQL injection
 - D. Fuzzing
17. Which of the following is an attack against servers hosting a directory service?
- A. XSS attack
 - B. LDAP injection attack
 - C. XSRF attack
 - D. Fuzzing attack

18. Your organization hosts a web site within a DMZ and the web site accesses a database server in the internal network. ACLs on firewalls prevent any connections to the database server except from the web server. Database fields holding customer data are encrypted and all data in transit between the web site server and the database server are encrypted. Which of the following represents the GREATEST risk to the data on the server?

- A. Theft of the database server
- B. XML injection
- C. SQL injection
- D. Sniffing

19. A security tester is sending random data to a program. What does this describe?

- A. Fuzzing
- B. Buffer overflow
- C. Integer overflow
- D. Command injection

20. Your organization is preparing to deploy a web-based application, which will accept user input. Which of the following will test the reliability of this application to maintain availability and data integrity?

- A. Secure coding
- B. Input validation
- C. Error handling
- D. Fuzzing

Chapter 7 Practice Question Answers

1. **D.** A distributed denial-of-service (DDoS) attack causes spikes in network traffic as multiple systems attempt to connect to a server and deplete the target’s resources. A smurf attack is an attack using directed broadcasts, and this might be a smurf attack if routers aren’t blocking directed broadcasts, but it could also be another type of DDoS attack. Flood guards protect against SYN flood attacks, and flood guards are not an attack method. A DoS attack comes from a single system.
2. **C.** Smurf attacks are blocked by preventing routers from passing directed broadcasts, especially border routers with direct access to the Internet. Flood guards protect against SYN (synchronize) flood attacks. Salting techniques add additional characters to passwords to thwart brute force attacks. Timestamps and sequence numbers are useful to protect against replay attacks, but not smurf attacks.
3. **B.** Timestamps and sequence numbers act as countermeasures against replay attacks. Blocking directed broadcasts prevents smurf attacks. Flood guards protect against SYN (synchronize) attacks. Salting protects against brute force attacks on passwords.

4. **A.** Account lockout policies help prevent brute force attacks by locking the account after an incorrect password or personal identification number (PIN) is entered too many times. This prevents someone from hacking into an account by guessing. A rainbow table is a type of attack. Salting passwords prevents some offline brute force attacks by adding characters to passwords before hashing them. Input validation prevents attacks such as buffer overflow and cross-site scripting, but wouldn't help here because an attacker guessing PINs is entering valid data.
5. **C.** A password salt is additional random characters added to a password before hashing the password, and it decreases the success of password attacks. Rainbow tables are used by attackers and contain precomputed hashes. Message digest 5 (MD5) is a hashing algorithm that creates hashes, but the scenario already states that passwords are hashed. Smurf is a type of attack using a directed broadcast and is not related to passwords.
6. **A.** A pharming attack attempts to redirect users from one web site to another web site. Although this is often done using DNS poisoning, it can also be done by rewriting the hosts file in a user's system. The 127.0.0.1 localhost entry is the default entry in the hosts file, and the second entry redirects the user to a different site. Whaling is a phishing attack that targets high-level executives. In session hijacking, an attacker records a user's credentials and uses them to impersonate the user. Phishing is the practice of sending email to users with the purpose of tricking them into revealing personal information (such as bank account information).
7. **B.** Typo squatting (or URL hijacking) uses a similar domain name to redirect traffic. In this scenario, the last two letters in CompTIA are swapped in the malicious domain name, and that site is attempting to download malware onto the user systems. A smurf attack is unrelated to web sites. Fuzzing tests an application's ability to handle random data. A replay attack attempts to replay data with the intent of impersonating one of the parties.
8. **B.** A zero-day attack takes advantage of an undocumented exploit or an exploit that is unknown to the public. A buffer overflow attack sends unexpected data to a system to access system memory or cause it to crash. Although some buffer overflow attacks are unknown, others are known. If the server isn't kept up to date with patches, it can be attacked with a known buffer overflow attack. Fuzzing sends random data to a system and can detect buffer overflows and zero-day attack methods, but the scenario doesn't indicate the user is sending random data. Session hijacking takes over a user's session and isn't related to an attack on a server.
9. **C.** Poor input validation often causes security vulnerabilities and can lead to major losses when exploited. Buffer overflow and cross-site request forgery (XSRF) are attacks that can be mitigated by input validation. They are not techniques used by developers. Hardening both operating systems and applications helps make them more secure from security vulnerabilities.
10. **C.** An attacker can use a web proxy to intercept the HTTP POST command. The attacker then modifies the data in the command and sends it to the web site. Command injection is a type of client-side injection attack that input validation thwarts. Flash cookies are used by Adobe Flash applets, but are not used to bypass input validation. Exception handling catches errors, allowing applications to handle them gracefully.
11. **B.** You should display a generic error message but log detailed information on the error. Detailed error messages to the user are often confusing to them and give attackers information they can use against the system. Logging generic information makes it more difficult to troubleshoot the problem later.
12. **A.** Buffer overflow attacks often cause an application to crash and expose system memory. Attackers then write malicious code into the exposed memory and use different techniques to get the system to run this code. None of the other attacks inserts malicious code into memory. Attackers attempt to embed HTML or JavaScript code in cross-site scripting (XSS) attacks, often to read cookies on a user's system. Extensible Markup Language (XML) injection attacks attempt to access or modify XML formatted data.
13. **B.** Buffer overflow attacks include a series of no operation (NOP) commands, such as hexadecimal 90 (x90). When successful, they can crash applications and expose memory, allowing attackers to run malicious code on the system. SQL injection attacks and Extensible Markup Language (XML) injection attacks do not use NOP commands. Zero-day attacks are unknown or undocumented, but attacks using NOP commands are known.
14. **D.** In a SQL injection attack, an attacker attempts to inject SQL commands into a query to access or manipulate data on a back-end database. A buffer overflow attack attempts to overload a system with too much data or unexpected data in an attempt to access system memory. A cross-site request forgery (XSRF) attack attempts to impersonate a user with HTML code. Fuzzing is a testing technique that sends random data into an application to see if the application can handle it.
15. **C.** A cross-site scripting (XSS) attack can be blocked by using input validation techniques to filter special characters such as the < and > characters used in HTML code. None of the other attackers requires the use of special characters. Sniffing captures data with a protocol analyzer. Spoofing hides the identity of the original entity. Pharming redirects a user from one web site to another web site.
16. **A.** A cross-site request forgery attack (XSRF) causes users to perform actions without their knowledge. This scenario indicates the user visited a web site, most likely through a malicious link, and the link initiated a purchase. None of the other attacks cause unsuspecting users to make purchases. A buffer overflow attacks a web site and attempts to access system memory. A SQL injection attack attempts to access data on a database server. Fuzzing sends random data to an application to test its ability to handle the random data.
17. **B.** A Lightweight Directory Application Protocol (LDAP) injection attack attempts to access data on servers hosting a directory service, such as a Microsoft domain controller hosting Active Directory. Cross-site scripting (XSS) and cross-site request forgery (XSRF) attacks attack web servers, not directory service servers. Fuzzing sends random data to see if the application can handle it, but it doesn't necessarily target servers hosting a directory service.
18. **C.** A SQL injection attack allows an attacker to send commands to the database server to access data. Encryption protects it on the server and in

transit, but the web server can decrypt it. Because the data in the database server is encrypted, theft of the server isn't a significant risk. There aren't any indications that the database server is replying with Extensible Markup Language (XML) data, so an XML injection attack isn't a risk. Because data is encrypted while in transit, sniffing isn't a significant risk.

19. **A.** Fuzz testing, or fuzzing, sends sending random data to an application with the purpose of testing the application's ability to handle the random data. In some cases, fuzzing can discover buffer overflow and integer overflow vulnerabilities, but just sending random data doesn't necessary cause buffer overflows or integer overflows. Command injection attacks send specific commands, not random data.

20. **D.** Fuzzing can test the application's ability to maintain availability and data integrity for some scenarios. Fuzzing sends random data to an application to verify the random data doesn't crash the application or expose the system to a data breach. Secure coding practices such as input validation and error- and exception-handling techniques protect applications, but do not test them.

Chapter 8

Managing Risk

CompTIA Security+ objectives covered in this chapter:

1.1 Implement security configuration parameters on network devices and other technologies.

- Protocol analyzers

2.1 Explain the importance of risk related concepts.

- Risk calculation (Likelihood, ALE, Impact, SLE, ARO, MTTR, MTTF, MTBF)
- Quantitative vs. qualitative, Vulnerabilities, Threat vectors, Probability / threat likelihood
- Risk-avoidance, transference, acceptance, mitigation, and deterrence

2.3 Given a scenario, implement appropriate risk mitigation strategies.

- User rights and permissions reviews
- Perform routine audits

2.8 Summarize risk management best practices.

- Risk assessment

3.2 Summarize various types of attacks.

- Malicious insider threat

3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques.

- Monitoring system logs (Event logs, Audit logs, Security logs, Access logs)

3.7 Given a scenario, use appropriate tools and techniques to discover security threats and vulnerabilities.

- Interpret results of security assessment tools
- Tools (Protocol analyzer, Vulnerability scanner, Port scanner, Passive vs. active tools, Banner grabbing)
- Risk calculations (Threat vs. likelihood)
- Assessment types (Risk, Threat, Vulnerability)
- Assessment technique (Baseline reporting, Code review, Determine attack surface, Review architecture, Review designs)

3.8 Explain the proper use of penetration testing versus vulnerability scanning.

- Penetration testing (Verify a threat exists, Bypass security controls, Actively test security controls, Exploiting vulnerabilities)
- Vulnerability scanning (Passively testing security controls, Identify vulnerability, Identify lack of security controls, Identify common misconfigurations, Intrusive vs. non-intrusive, Credentialed vs. non-credentialed, False positive)
- Black box, White box, Gray box

5.3 Install and configure security controls when performing account management, based on best practices.

- User access reviews, Continuous monitoring

**

As a security professional, you need to be aware of the different security issues associated with threats, vulnerabilities, and risks, and the tools available to combat them. This chapter digs into risk management concepts, including risk assessment methods, and methods used to check for vulnerabilities. It also covers some specific tools such as protocol analyzers and port scanners.

Identifying Risk

Risk is the likelihood that a threat will exploit a vulnerability. A *vulnerability* is a weakness, and a *threat* is a potential danger. The result is a negative impact on the organization. *Impact* refers to the magnitude of harm that can be caused if a threat exercises a vulnerability.

For example, a system without up-to-date antivirus software is vulnerable to malware. Malware written by malicious attackers is the threat. The likelihood that the malware will reach a vulnerable system represents the risk. Depending on what the malware does, the impact may be an unbootable computer, loss of data, or a remote-controlled computer that has joined a botnet.

However, the likelihood of a risk occurring isn't 100 percent. An isolated system without Internet access, network connectivity, or USB ports has a very low likelihood of malware infection. The likelihood significantly increases for an Internet-connected system, and it increases even more if a user visits risky web sites and downloads and installs unverified files.

It's important to realize that you can't eliminate risk. Sure, you can avoid information technology (IT) risks completely by unplugging your computer and burying it. However, that's not very useful. Instead, users and organizations practice risk management to reduce the risks.

You probably practice risk management every day. Driving or walking down roads and streets can be a very dangerous activity. Car-sized bullets are speeding back and forth, representing significant risks to anyone else on the road. However, you mitigate these risks with caution and vigilance. The same occurs with computers and networks. An organization mitigates risks using different types of security controls.

Threats and Threat Vectors

A threat is a potential danger. Within the realm of CompTIA Security+, a *threat* is any circumstance or event that can compromise the confidentiality, integrity, or availability of data or a system. A *threat vector* (also called an attack vector) refers to the method used to activate the threat and can originate from three primary sources: external (outsiders), internal (insiders), and the supply chain (suppliers).

Types of Threats

Threats come in different forms, including the following:

- **Natural threats.** This could include hurricanes, floods, tornadoes, earthquakes, landslides, electrical storms, and other similar events. On a less drastic scale, a natural threat could also mean hardware failure.
- **Malicious human threats.** Attackers regularly launch different types of attacks, including network attacks, system attacks, and the release of malware.
- **Accidental human threats.** Users can accidentally delete or corrupt data, or accidentally access data that they shouldn't be able to access. Even administrators can unintentionally cause system outages. The common cause is by a well-meaning administrator making a configuration change to fix one problem but inadvertently causing another one.
- **Environmental threats.** This includes long-term power failure, which could lead to chemical spills, pollution, or other possible threats to the environment.

Different locations have different threats. When evaluating threats, it's important to consider the likelihood of the threat. For example, I live in Virginia Beach, Virginia, and while we're concerned about the natural threat of hurricanes during the hurricane season, we aren't very concerned about earthquakes. My sister is a business continuity expert and she lives in San Francisco and works in Silicon Valley. She helps companies prepare for risks associated with earthquakes there, but she spends very little time or energy considering the risk of a hurricane hitting San Francisco.

Malicious Insider Threat

A *malicious insider* is anyone who has legitimate access to an organization's internal resources, but exploits this access for personal gain or damage against the organization. This person's actions can compromise confidentiality, integrity, and availability of the organization's assets.

Malicious insiders have a diverse set of motivations. For example, some malicious insiders are driven by greed and simply want to enhance their finances, while others want to exact revenge on the organization. They may steal files that include valuable data, install or run malicious scripts, redirect funds to their personal accounts, or take any of countless other actions.

Most employees are overwhelmingly honest, but a single malicious insider can launch a successful attack and cause significant damage to the company. Because of this, most organizations implement basic controls to prevent potential problems.

For example, Chapter 2, "Exploring Control Types and Methods," discusses the principle of least privilege. This principle ensures that employees have only the rights and permissions to perform their assigned tasks and functions, and limits the potential damage they can cause if they become malicious insiders. Chapter 11, "Exploring Operational Security," discusses other policies, such as job rotation, separation of duties, and mandatory vacations.

Threat Assessments

Threat assessments help an organization identify and categorize threats. They attempt to predict the threats against a system or application along with the likelihood a threat vector will activate the threat. Threat assessments also attempt to identify the potential impact from these threats. Once the organization identifies and prioritizes threats, it identifies security controls to protect against the most serious threats.

Organizations have limited resources, so it's not possible to protect against all threats. However, threat assessments improve the security posture of any system or application by ensuring that the resources aren't squandered on low-priority threats.

Vulnerabilities

A *vulnerability* is a flaw or weakness in software or hardware, or a weakness in a process that a threat could exploit resulting in a security breach.

Examples of vulnerabilities include:

- **Lack of updates.** If systems aren't kept up to date with patches, hotfixes, and service packs, they are vulnerable to bugs and flaws in the software.
- **Default configurations.** Hardening a system includes changing systems from their default hardware and software configurations, including changing default usernames and passwords. If systems aren't hardened, they are more susceptible to attacks. Chapter 5, "Securing Hosts and Data," covers hardening systems in more depth.
- **Lack of malware protection or updated definitions.** Antivirus and anti-spyware methods protect systems from malware, but if they aren't used and kept up to date, systems are vulnerable to malware attacks. Chapter 6, "Understanding Malware and Social Engineering," covers malware types and methods used to protect systems from malware attacks.
- **Lack of firewalls.** If personal and network firewalls aren't enabled or configured properly, systems are more vulnerable to network and Internet-based attacks.
- **Lack of organizational policies.** If job separation, mandatory vacations, and job rotation policies aren't implemented, an organization may be more susceptible to fraud and collusion from employees.

Not all vulnerabilities are exploited. For example, a user may install a wireless router using the defaults. It is highly vulnerable to an attack, but that doesn't mean that an attacker will discover it and attack. In other words, just because the wireless router has never been attacked, it doesn't mean that it isn't vulnerable. At any moment, a war driving attacker can drive by and exploit the vulnerability.

Risk Management

Risk management is the practice of identifying, monitoring, and limiting risks to a manageable level. It doesn't eliminate risks, but instead identifies methods to limit or mitigate them. The amount of risk that remains after managing risk is *residual risk*.

The primary goal of risk management is to reduce risk to a level that the organization will accept. Senior management is ultimately responsible for residual risk—the amount of risk that remains after mitigating risk. Management must choose a level of acceptable risk based on their organizational goals. They decide what resources (such as money, hardware, and time) to dedicate to mitigate the risk.

There are multiple risk management methods available to an organization. They include:

- **Risk avoidance.** An organization can avoid a risk by not providing a service or not participating in a risky activity. For example, an organization may evaluate an application that requires multiple open ports on the firewall that it considers too risky. It can avoid the risk by purchasing another application.
- **Risk transference.** The organization transfers the risk to another entity, or at least shares the risk with another entity. The most common method is by purchasing insurance. Another method is by outsourcing, or contracting a third party.
- **Risk acceptance.** When the cost of a control outweighs a risk, an organization will often accept the risk. For example, spending \$100 in hardware locks to secure a \$15 mouse doesn't make sense. Instead, the organization accepts the risk of someone stealing the mouse. Similarly, even after implementing controls, residual risk remains and the organization accepts this residual risk.
- **Risk mitigation.** The organization implements controls to reduce the risk. These controls may reduce the vulnerabilities or reduce the impact of the threat. For example, up-to-date antivirus software mitigates the risks of malware.
- **Risk deterrence.** An organization can deter a risk by implementing some security controls. For example, a security guard can deter an attacker from trying to access a secure area. Similarly, cameras can mitigate risks associated with theft.

Some security professionals identify the first four methods of risk management but don't include risk deterrence. Instead, they include deterrence methods within the risk mitigation category. However, the CompTIA Security+ objectives list these five.

Remember this

It is not possible to eliminate risk, but you can take steps to manage it. An organization can avoid a risk by not providing a service or not participating in a risky activity. Insurance transfers the risk to another entity. You can mitigate risk by implementing controls, but when the cost of the controls exceeds the cost of the risk, an organization accepts the remaining, or residual risk.

Risk Assessment

A risk assessment, or risk analysis, is an important task in risk management. It quantifies or qualifies risks based on different values or judgments. A risk assessment starts by first identifying assets and asset values. This helps an organization focus on the high-value assets and avoid wasting time on low-value assets.

It then identifies threats and vulnerabilities and determines the likelihood a threat will attempt to exploit a vulnerability. A risk assessment attempts to identify the impact of potential threats and identify the potential harm, and prioritizes risks based on the likelihood and impact. Last, a risk assessment includes recommendations on what controls to implement to mitigate risks.

A risk assessment is a point-in-time assessment, or a snapshot. In other words, it assesses the risks based on current conditions, such as current threats, vulnerabilities, and existing controls. For example, consider a library computer that has up-to-date antivirus protection and cannot access the Internet. Based on these conditions, the risks are low. However, if administrators connect the system to the Internet, or fail to keep the antivirus software up to date, the risk increases.

It's common to perform risk assessments on new systems or applications. For example, if an organization is considering adding a new service or application that can increase revenue, it will often perform a risk assessment. This helps it determine if the potential risks may offset the potential gains.

Risk assessments use quantitative measurements or qualitative measurements. Quantitative measurements use numbers, such as a monetary figure representing cost and asset values. Qualitative measurements use judgments. Both methods have the same core goal of helping management make educated decisions based on priorities.

Quantitative Risk Assessment

A *quantitative risk assessment* measures the risk using a specific monetary amount. This monetary amount makes it easier to prioritize risks. For example, a risk with a potential loss of \$30,000 is much more important than a risk with a potential loss of \$1,000.

The asset value is an important element in a quantitative risk assessment. It may include the revenue value or replacement value of an asset. A web server may generate \$10,000 in revenue per hour. If the web server fails, the company will lose \$10,000 in direct sales each hour it's down, plus the cost to repair it. It can also result in the loss of future business if customers take their business elsewhere. In contrast, the failure of a library workstation

may cost a maximum of \$1,000 to replace it.

One quantitative model uses the following values to determine risks:

- **Single loss expectancy (SLE).** The SLE is the cost of any single loss.
- **Annual rate of occurrence (ARO).** The ARO indicates how many times the loss will occur in a year. If the ARO is less than 1, the ARO is represented as a percentage. For example, if you anticipate the occurrence once every two years, the ARO is 50 percent or .5.
- **Annual loss expectancy (ALE).** The ALE is the $SLE \times ARO$.

Imagine that employees at your company lose, on average, one laptop a month. Thieves have stolen them when employees left them in conference rooms during lunch, while they were on location at customer locations, and from training rooms.

Someone suggested purchasing hardware locks to secure these laptops for a total of \$1,000. These locks work similar to bicycle locks and allow employees to wrap the cable around a piece of furniture and connect into the laptop. A thief needs to either destroy the laptop to remove the lock or take the furniture with them when stealing the laptop. Should your company purchase them? With a little analysis, the decision is easy.

You have identified the average cost of these laptops, including the hardware, software, and data, as \$2,000 each. This assumes employees do not store entire databases of customer information or other sensitive data on the systems, which can easily result in much higher costs. You can now calculate the SLE, ARO, and ALE as follows:

- **SLE.** The value of each laptop is \$2,000, so the SLE is \$2,000.
- **ARO.** Employees lose about one laptop a month, so the ARO is 12.
- **ALE.** You calculate the ALE as $SLE \times ARO$, so $\$2,000 \times 12 = \$24,000$.

Security experts estimate that these locks will reduce the number of lost or stolen laptops from 12 a year to only 2 a year. This changes the ALE from \$24,000 to only \$4,000 (saving \$20,000 a year). In other words, the organization can spend \$1,000 to save \$20,000. It doesn't take a rocket scientist to see that this is a good fiscal decision, saving a net of \$19,000. Buy them.

Managers use these two simple guidelines for most of these decisions:

- If the cost of the control is less than the savings, purchase it.
- If the cost of the control is greater than the savings, accept the risk.

The organization might be considering other controls, such as a combination of hardware locks, biometric authentication, LoJack for Laptops, and more. The final cost of all of these controls is \$30,000 per year. Even if a laptop is never stolen again, the company is spending \$30,000 to save \$24,000, resulting in a higher net loss—they're losing \$6,000 more a year.

Admittedly, a company could choose to factor in other values, such as the sensitivity of data on the laptops, and make a judgment to purchase these controls. However, if they're using a quantitative risk assessment, these values would need to be expressed in monetary terms.

Although you would normally know the SLE and ARO and use these to calculate the ALE, you might occasionally have the SLE and ALE, but not know the ARO. Using basic algebra, you can reformat the formula. Any of these are valid:

- $ALE = SLE \times ARO$
- $ARO = ALE / SLE$
- $SLE = ALE / ARO$

Remember this

A quantitative risk assessment uses specific monetary amounts to identify cost and asset values. The SLE identifies the amount of each loss, the ARO identifies the number of failures in a year, and the ALE identifies the expected annual loss. You calculate the ALE as $SLE \times ARO$. A qualitative risk assessment uses judgment to categorize risks based on probability and impact.

Qualitative Risk Assessment

A *qualitative risk assessment* uses judgment to categorize risks based on probability and impact. *Probability* is the likelihood that an event will occur, such as the likelihood that a threat will attempt to exploit a vulnerability. *Impact* is the negative result of the event, such as loss of confidentiality, integrity, or availability of a system or data.

Notice that this is much different from the exact numbers provided by a quantitative assessment that uses monetary figures. You can think of quantitative as using a quantity or a number, whereas qualitative is related to quality, which is often a matter of judgment.

Some qualitative risk assessments use surveys or focus groups. They canvass experts to provide their best judgments and then tabulate the results. For example, a survey may ask the experts to rate the probability and impact of risks associated with a web server selling products on the Internet and a library workstation without Internet access. The experts would use words such as *low*, *medium*, and *high* to rate them.

They could rate the probability of a web server being attacked as high, and if the attack takes the web server out of service, the impact is also high. On the other hand, the probability of a library workstation being attacked is low, and, even though a library patron may be inconvenienced, the impact is also low.

It's common to assign numbers to these judgments. For example, you can use terms such as low, medium, and high, and assign values of 1, 5, and

10, respectively. The experts assign a probability and impact of each risk using low, medium, and high, and when tabulating the results, you change the words to numbers. This makes it a little easier to calculate the results.

In the web server and library computer examples, you can calculate the risk by multiplying the probability and the impact:

- **Web server.** High probability and high impact: $10 \times 10 = 100$.
- **Library computer.** Low probability and low impact: $1 \times 1 = 1$.

Management can look at these numbers and easily determine how to allocate resources to protect against the risks. They would allocate more resources to protect the web server than the library computer.

One of the challenges with a qualitative risk assessment is gaining consensus on the probability and impact. Unlike monetary values that you can validate with facts, probability and impact are often subject to debate.

Documenting the Assessment

The final phase of the risk assessment is the report. This identifies the risks discovered during the assessment and the recommended controls. As a simple example, a risk assessment on a database-enabled web application may discover that it's susceptible to SQL injection attacks. The risk assessment will then recommend rewriting the web application with input validation techniques or stored procedures to protect the database.

Management uses this to decide which controls to implement and which controls to accept. In many cases, a final report documents the managerial decisions. Of course, management can decide not to implement a control, but instead accept a risk.

Think how valuable this report will be for an attacker. They won't need to dig to identify vulnerabilities or controls. Instead, the report lists all the details. Even when management approves controls to correct the vulnerabilities, it may take some time to implement them. Because of this, the results of a risk assessment are highly protected. Normally, only executive management and security professionals will have access to these reports.

Using Metrics to Identify Risk

When identifying risks associated with hardware, security experts and administrators use several metrics. These include:

- **Mean time between failures (MTBF).** The mean time between failures (MTBF) provides a measure of a system's reliability and is usually represented in hours. More specifically, the MTBF identifies the average (the arithmetic mean) time between failures. Higher MTBF numbers indicate a higher reliability of a product or system. Administrators and security experts attempt to identify the MTBF for critical systems with a goal of predicting potential outages.
- **Mean time to failure (MTTF).** The mean time to failure (MTTF) is the length of time you can expect a device to remain in operation before it fails. It is similar to MTBF, but the primary difference is that the MTBF metric indicates you can repair the device after it fails. The MTTF metric indicates that you will not be able to repair a device after it fails.
- **Mean time to recover (MTTR).** The mean time to recover (MTTR) identifies the average (the arithmetic mean) time it takes to restore a failed system. In some cases, people interpret MTTR as the mean time to repair, and both mean essentially the same thing. Organizations that have maintenance contracts often specify the MTTR as a part of the contract. The supplier agrees that it will, on average, restore a failed system within the MTTR time. The MTTR does not provide a guarantee that it will restore the system within the MTTR every time. Sometimes it may take a little longer and sometimes it may be a little quicker, with the average defined by the MTTR.

Checking for Vulnerabilities

Vulnerabilities are weaknesses, and by reducing vulnerabilities, you can reduce risks. That sounds simple enough. However, how do you identify the vulnerabilities that present the greatest risks? Common methods are with vulnerability assessments, vulnerability scans, and penetration tests.

To understand how these are used, it's worthwhile to understand how attackers may look for targets. The following section outlines common attack methods, followed by some details on vulnerability assessments, vulnerability scans, and penetration tests.

Anatomy of an Attack

From a defensive perspective, it's valuable to understand how attackers operate. Many penetration testers use similar methodologies, so you can also apply this knowledge to penetration testing. Although there is no single definition that identifies all attackers, one thing is clear: Attackers are sophisticated and clever. They should not be underestimated.

Imagine an organization that has unlimited funds to launch attacks. They could be government employees employed by another country, or they could be a group of criminals. They may be trying to steal secrets or steal some of the millions of dollars that attackers are pocketing from businesses and individuals monthly. How will they go about it?

They often combine reconnaissance with fingerprinting techniques to identify targets. Reconnaissance provides a big-picture view of a network, including the Internet Protocol (IP) addresses of a target network. Fingerprinting then homes in on individual systems to provide details of each. Some of the attackers may be experts on reconnaissance, while others are experts on different elements of fingerprinting. Once they identify their targets, attackers use this information to launch an attack. The following sections identify one possibility of how these attackers may operate.

Identifying IP Addresses of Targets

One group in this organization identifies IP addresses of live systems as potential targets. For example, they may be interested in government or military systems in a specific area, or a certain company operating in a specific city.

IP addresses are assigned geographically, so this is a little easier than it may seem. You've probably been surfing the Internet and seen advertisements for your city, even though the web page isn't local. The advertisement identifies your location based on your IP address and targets the ad. Attackers have access to the same information as the advertisers. As an example, look at the IP address locator at <http://www.whatismyip.com/>. It defaults to your public IP address, but you can plug in any IP address to identify its location.

Once attackers identify a geographical range of IP addresses, attackers can use an Internet Control Message Protocol (ICMP) sweep or host enumeration sweep to identify systems that are operational in that range. This is similar to sending a ping to each IP address in that range, and tools that perform ICMP sweeps are commonly called ping scanners. As mentioned in Chapter 3, "Understanding Basic Network Security," it's possible to block ICMP at firewalls to reduce the success of ICMP sweeps.

Identifying Open Ports with a Port Scanner

The next group in the organization takes the list of IP addresses and identifies open ports on each system. Chapter 3 covers ports in greater depth. As a reminder, many protocols use well-known ports. For example, Telnet uses port 23 and Hypertext Transfer Protocol (HTTP) uses port 80. If a port scanner detects port 80 is open, it's likely that the HTTP protocol is running and it may be a web server.

Additionally, ports also identify applications running on a system. For example, peer-to-peer (P2P) software, used for file sharing, uses specific ports for communications with other peers on the Internet. A port scan can discover these ports and detect P2P software running on a system.

A common Transmission Control Protocol (TCP) port scan sends a TCP SYN (synchronize) packet to a specific port of a server as part of the TCP three-way handshake. If the server responds with a SYN/ACK (synchronize/acknowledge) packet, the scanner knows the port is open. However, instead of completing the three-way handshake, the scanner can send an RST (reset) packet to reset the connection and then repeat the process with a different port.

Even though it's not recommended for services to use ports other than the well-known ports for specific services, it is possible. Because of this, an open port doesn't definitively say the related service or protocol is running.

Some port scanners send additional queries to the system's open ports. For example, if port 25 is open, indicating it's running Simple Mail Transfer Protocol (SMTP), the scanner can send SMTP queries to the system and analyze the response. These queries provide verification that the protocol is running and often include additional details on the system. Similarly, HTTP queries can identify if it's a Windows Internet Information Services (IIS) web server or an Apache web server running on a Linux system.

Each open port represents a potential attack vector, so by identifying open ports, attackers can determine the attack surface. As a reminder, you can reduce the attack surface by disabling all unused services and removing all unneeded protocols, as discussed in Chapter 5. This is a key step in hardening a system, or making it more secure from the default configuration.

Many vulnerability-scanning tools like Nmap, Netcat, and Nessus include port-scanning abilities and scan systems to determine open ports. Security professionals use these tools for vulnerability scans within their networks, but attackers also use these same tools. In addition to scanning for open ports, these tools can also fingerprint a system.

Remember this

A port scanner can help determine what services and protocols are running on a remote system by identifying open ports. Port scanners typically take additional steps to verify the port is open.

Fingerprint System

Attackers also attempt to fingerprint the system to identify additional details. Operating system fingerprinting identifies the operating system. For example, is this a Linux system or a Windows system? A fingerprinting attack sends specific protocol queries to the server and analyzes the responses. These responses can verify that a service is running and often include other details about the operating system because different operating systems often respond differently to specific queries.

Chapter 7, “Identifying Advanced Attacks,” introduces the Xmas attack as a specific type of port scan. It analyzes the returned packets to identify the operating system of the target system, and sometimes even the version of the operating system, in addition to identifying open ports. Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) can reduce the success of port scanning and fingerprinting scans.

Banner Grabbing

Banner grabbing is a technique used to gain information about a remote server and is often used as part of a fingerprinting attack. Many attackers use Telnet because it’s relatively simple to use. For example, the following command attempts to connect to the gcgapremium.com server using the HTTP protocol on port 80: **telnet gcgapremium.com 80**. If successful, the server returns a Hypertext Markup Language (HTML) banner providing information on the server. The banner might look something like the following:

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
```

```
<html><head><title>501 Method Not Implemented</title>
```

```
</head><body>
```

```
<h1>Method Not Implemented</h1>
```

```
<p>GET to /index.html not supported.<br /></p>
```

```
<p>Additionally, a 404 Not Found error was encountered.</p><hr>
```

```
<address>Apache/2.2.25 (Unix) mod_ssl/2.2.25 OpenSSL/1.0.0-fips mod_auth_passthrough/2.1 mod_bwlimited/1.4 FrontPage/5.0.2.2635 Server  
at 72.52.230.233 Port 80</address>
```

```
</body></html>
```

Most of this is formatting. However, the information in the address section provides a lot of information on the web server. It shows this is a Unix server running the Apache web server software along with additional information. If you want to see how to do this, check out the Banner Grabbing Lab in the online exercises for this book at <http://gcgapremium.com/labs/>.

Identifying Vulnerabilities

At this stage, attackers know the IP address of live systems, what operating systems they’re running, and what protocols they’re running. This information is passed to the appropriate experts. For example, some attackers may be experts at attacking IIS web servers and Windows systems, while others are experts at attacking Apache web servers and Linux or Unix systems.

If it’s a web server, attackers may check to see if input validation techniques are in place. If not, it may be susceptible to buffer overflow, command injection, SQL injection, or cross-site scripting attacks. If it’s an application server with known default accounts and passwords, the attacker checks to see if the defaults are still available. Chapter 7 presented information on each of these attacks.

Many vulnerability scanners can easily check for current patches. If systems aren’t kept up to date with current patches, they are susceptible to known vulnerabilities. Again, vulnerability scanners can identify vulnerabilities, and both security professionals and attackers can use them.

Attack

Attackers now have a list of systems and their vulnerabilities. They may have attack tools they can run immediately to exploit vulnerable systems. Other times, experts may write code to exploit a new vulnerability. However, you can bet that they will attack.

The reconnaissance and fingerprinting stages may take days, weeks, or months. Some attacks are quick after some detailed planning, whereas other attacks linger as long as the attacker is undetected or until the attacker has completed the mission.

For example, when attackers extract or exfiltrate data, they often get in and get out as soon as possible. This often requires them to use privilege escalation tactics to gain elevated rights and permissions to access the data. Escalating their privileges is likely to sound an alarm. Once they’re detected, they can expect to be blocked, so they must get as much data as possible, as quickly as possible. However, sometimes the attacks are undetected, and they can continue the attacks for days.

It’s also worth pointing out that attackers often launch attacks through other systems. They take control of remote computers through different types of malware and launch the attacks through these systems. It is difficult, though not impossible, to track these attacks back to the actual source.

As the last step in the attack, many attackers attempt to erase or modify the logs. The goal is to remove traces of their attack.

Putting It All Together

Is it possible that governments have dedicated teams working together to identify exploitable vulnerabilities? Absolutely. Is it possible that criminals have the ability to organize their efforts to identify targets of opportunity? Count on it. An advanced persistent threat (APT) is a group that has both the capability and intent to launch sophisticated and targeted attacks, and there is a lot of evidence that they exist and are active.

Information security company Mandiant released a report in 2013 concluding that at least one APT operating from China is likely government-sponsored. You can find the report by Googling “Mandiant APT1.” Mandiant concluded that the group they named APT1 operates as Unit 61398 of the People’s Liberation Army (PLA) inside China. Mandiant estimates that APT1 includes over 1,000 servers and between dozens and hundreds of individual operators and has:

- Released at least 40 different families of malware
- Stolen hundreds of terabytes of data from at least 141 organizations
- Maintained access to some victim networks for over four years before being detected
- Established footholds within many networks after email recipients opened malicious files that installed backdoors, allowing attackers remote access

Chinese officials have denied these claims.

Cyberwarfare and cybercrime are similar to spying and espionage. Espionage is the process of gathering multiple innocuous details that form the individual pieces of a much larger picture. Just as the pieces of a jigsaw puzzle eventually come together to complete a picture, the individual details of any system come together to fingerprint it. Security professionals need to be aggressive at closing all the holes to limit the amount of information that is available to any attacker. It’s also become increasingly more important to train users on how criminals deliver malware through spam and phishing emails.

Of course, some criminals don’t have large organizations. A handful of attackers can combine their skills to identify a niche or specialty. They could become experts at SQL injection attacks and only need to look for web sites without input validation. They could exploit a specific vulnerability to manage a botnet. The possibilities are endless.

However, security professionals can’t afford to protect against niches only. Security professionals must protect against all attacks.

Vulnerability Assessment

The overall goal of a vulnerability assessment is to assess the security posture of systems and networks. They identify vulnerabilities, or weaknesses, within systems, networks, and organizations, and are part of an overall risk management plan.

Vulnerability assessments can include information from a wide variety of sources. This includes reviewing security policies and logs, interviewing personnel, and testing systems. Assessments often use vulnerability scans and penetration tests, covered in more depth later in this chapter. A vulnerability assessment often includes the following high-level steps:

- Identify assets and capabilities.
- Prioritize assets based on value.
- Identify vulnerabilities and prioritize them.
- Recommend controls to mitigate serious vulnerabilities.

Many organizations perform vulnerability assessments internally. Organizations also occasionally hire external security professionals to complete external assessments.

Remember this

A vulnerability assessment determines the security posture of a system or network by identifying vulnerabilities and weaknesses. Assessments include the results of various tools, such as vulnerability scanners, audits, and reviews.

Vulnerability Scanning

A key part of a vulnerability assessment is a vulnerability scan. Security administrators use vulnerability scanners as a management control to identify which systems are susceptible to attacks. Vulnerability scanners identify a wide range of weaknesses and known security issues that attackers can exploit. Most vulnerability scanners combine multiple features into a single package. A vulnerability scanner includes the following capabilities:

- Identifying vulnerabilities
- Identifying misconfigurations
- Passively testing security controls
- Identifying lack of security controls

Identifying Vulnerabilities and Misconfigurations

Vulnerability scanners utilize a database or dictionary of known vulnerabilities and test systems against this database. For example, the MITRE Corporation maintains the Common Vulnerabilities and Exposures (CVE) list, which is a dictionary of publicly known security vulnerabilities and exposures. This is similar to how antivirus software detects malware using virus signatures. The difference is that the CVE is one public list funded by the U.S. government, whereas antivirus vendors maintain proprietary signature files.

Additionally, attackers often look for systems that are misconfigured, but vulnerability scanners can detect some common misconfiguration settings. Some of the vulnerabilities and common misconfigurations discovered by a vulnerability scanner include:

- **Open ports.** Open ports can signal a vulnerability, especially if administrators aren't actively managing the services associated with these ports. For example, not all web servers use File Transfer Protocol (FTP) so if TCP ports 20 and 21 are open, it indicates a potential vulnerability related to FTP. Similarly, Telnet uses port 23 but Secure Shell (SSH) using port 22 is a recommended alternative.
- **Weak passwords.** Many scanners include a password cracker that can discover weak passwords or verify that users are creating strong passwords in compliance with an organization's policy. It is more efficient to use a technical password policy to require and enforce the use of strong passwords. However, if this isn't possible, administrators use a separate password cracker to discover weak passwords.
- **Default accounts and passwords.** Operating systems and applications can have default usernames and passwords. Basic operating system and application hardening steps should remove the defaults, and a scan can discover the weaknesses if operating systems and applications aren't hardened. For example, some SQL database systems allow the sa (system administrator) account to be enabled with a blank password. Scanners such as Nessus will detect this.
- **Sensitive data.** Some scanners include data loss prevention (DLP) techniques to detect sensitive data sent over the network. For example, a DLP system can scan data looking for patterns such as Social Security numbers or key words that identify classified or proprietary data.
- **Security and configuration errors.** Vulnerability scans can also check the system against a configuration or security baseline to identify unauthorized changes.

Administrators can scan specific systems or an entire network. For example, many organizations perform periodic scans on the entire network to detect vulnerabilities. If an administrator makes an unauthorized change resulting in a vulnerability, the scan can detect it. Similarly, if a rebuilt system is missing some key security settings, the scan will detect them. It's also possible to scan a new system before or right after it's deployed.

Passively Testing Security Controls

An important point about a vulnerability scan is that it does not attempt to exploit any vulnerabilities. Instead, a vulnerability scan is a passive attempt to identify weaknesses. This ensures that the testing does not interfere with normal operations. Security administrators then assess the vulnerabilities to determine which ones to mitigate. In contrast, a penetration test (covered later in this chapter) is an active test that attempts to exploit vulnerabilities.

Identifying Lack of Security Controls

Vulnerability scanners can also identify missing security controls, such as the lack of up-to-date patches or the lack of antivirus software. Although many patch management tools include the ability to verify systems are up to date with current patches, vulnerability scanners provide an additional check to detect unpatched systems.

Remember this

A vulnerability scanner can identify vulnerabilities, misconfigured systems, and the lack of security controls such as up-to-date patches. Vulnerability scans are passive and have little impact on a system during a test. In contrast, a penetration test is intrusive and can potentially compromise a system.

False Positive

Unfortunately, vulnerability scanners aren't perfect. Occasionally, they report a vulnerability when it doesn't actually exist. In other words, the scan reports a positive on a known vulnerability, but the report is false. As an example, a vulnerability scan on a server might report that the server is missing patches related to a database application, but the server doesn't have a database application installed.

This is similar to false positives in an intrusion detection system (IDS) where the IDS alerts on an event, but the event isn't an actual intrusion. Similarly, an antivirus scanner can identify a useful application as malware, even though the application does not have any malicious code. False positives can result in higher administrative overhead because administrators have to investigate them.

Remember this

A false positive from a vulnerability scan indicates the scan detected a vulnerability, but the vulnerability doesn't actually exist.

Other Assessment Techniques

Vulnerability assessments are broader than just vulnerability scans. A vulnerability scanner can discover technical vulnerabilities, but an organization can also have nontechnical vulnerabilities that go beyond technical controls.

A vulnerability assessment can also check for nontechnical vulnerabilities. For example, Chapter 6 mentioned tailgating, where an employee follows closely behind another employee without using credentials. One employee uses a proximity card to open a door and other employees follow. If employees are tailgating, can an attacker do the same? Theoretically, yes, but in many cases, management wants more than theory. A vulnerability assessment can include a test to see if a visitor can access secure spaces without credentials.

Similarly, employees may be susceptible to social engineering attacks. An attacker may use low-tech methods to trick employees into revealing sensitive information. Educated employees often recognize these techniques, but even if a company provides training, it doesn't necessarily mean that the employees are educated. A vulnerability assessment can verify what training was effective and sometimes identify which employees represent the highest risks.

For example, users should not give their password out to anyone, and many organizations regularly remind users of this security practice. However, will users give out their password?

I remember one vulnerability assessment performed within a bank. The testers drafted an official-looking email explaining a fictitious problem, but linked it to an actual internal server migration. The email indicated that due to the migration, there was a problem with the accounts and users would lose account access unless they provided their password. All of the employees attended training less than a month earlier on the importance of not giving out their passwords. Still, over 35 percent of the employees provided their password in response to this email.

Other assessment techniques include:

- **Baseline reporting.** Chapter 5 presented information on baselines, including security baselines and configuration baselines. Many vulnerability assessment tools can perform baseline reviews by comparing current security and configuration data with a baseline to detect changes. These changes can introduce vulnerabilities, so they should be investigated.
- **Code review.** Chapter 7 discussed methods of ensuring that code is secure and mentioned the importance of code review. A code review goes line-by-line through the code and can help detect vulnerabilities, such as race conditions or susceptibility to buffer overflow attacks. Other programmers often perform them as a peer assessment.
- **Attack surface review.** The attack surface refers to the attack vectors available on a system, such as open ports. By hardening a system, you

reduce the attack surface and a vulnerability assessment evaluates a system to determine if it is adequately hardened.

- **Architecture review.** Security experts review the network architecture by examining the network and looking for potential vulnerabilities. For example, a review may discover that a database server is located within a demilitarized zone (DMZ) and is accessible from the Internet. The review can recommend moving the database server behind an additional firewall.
- **Design review.** Security experts can also identify vulnerabilities by reviewing designs. This includes reviewing the physical layout of a building, the layout of the network, or how an application interacts with other applications or systems. Security is easier to implement early in the design stage than it is to implement later, and a design review ensures that systems and software are developed properly.

Remember this

A baseline review identifies changes from the standard configuration. Code reviews review software line-by-line to identify potential vulnerabilities such as race conditions or susceptibility to buffer overflow attacks. Design reviews ensure that systems and software are developed properly.

Credentialed Versus Noncredentialed

Vulnerability scanners can run as a credentialed scan using the credentials of an account, or a noncredentialed scan without any user credentials.

Attackers typically do not have the credentials of an internal account so when they run scans against systems, they run noncredentialed scans.

Security administrators often run credentialed scans with the privileges of an administrator account. This allows the scan to check security issues at a much deeper level than a noncredentialed scan. Additionally, because the credentialed scan has easier access to internal workings of systems, it results in a lower impact on the tested systems, along with more accurate test results.

It's worth mentioning that attackers typically start without any credentials but using privilege escalation techniques, they often gain administrative access. This allows them to run a credentialed scan against a network if desired. Similarly, even though a credentialed scan is typically more accurate, administrators often run noncredentialed scans to see what an attacker without credentials would see.

Penetration Testing

A *penetration test* (sometimes called a pentest) actively assesses deployed security controls within a system or network. It starts with a vulnerability scan but takes it a step further and actually tries to exploit the vulnerability by simulating or performing an attack.

Security testers typically perform a penetration test to demonstrate the actual security vulnerabilities within a system. This can help the organization determine the impact of a threat against a system. In other words, it helps an organization determine the extent of damage that an attacker could inflict by exploiting a vulnerability.

Although it's not as common, it's also possible to perform a penetration test to determine how an organization will respond to a compromised system. This allows an organization to demonstrate security vulnerabilities and flaws in policy implementation. For example, many organizations may have perfect policies on paper. However, if employees aren't consistently following the policies, a penetration test can accurately demonstrate the flaws.

Many penetration tests include the following activities:

- Verify a threat exists.
- Bypass security controls.
- Actively test security controls.
- Exploit vulnerabilities.

For example, an organization could hire an external tester to test the security of a web application. A first step could check for a SQL injection vulnerability. If the vulnerability exists, the tester could then launch a SQL injection attack to harvest user credentials from a database. The attacker can then use these credentials to exploit other areas of the system. If the database included credentials of elevated accounts, the attacker can use these for privilege escalation and exploit other system vulnerabilities.

Because a penetration test can exploit vulnerabilities, it has the potential to disrupt actual operations and cause system instability. Because of this, it's important to define boundaries strictly for a test. Ideally, the penetration test will stop right before performing an exploit that can cause damage or result in an outage. However, some tests cause unexpected results.

Testers sometimes perform penetration tests on test systems rather than the live production systems. For example, an organization may be hosting a web application accessible on the Internet. Instead of performing the test on the live server and affecting customers, penetration testers or administrators configure another server with the same web application. If a penetration test cripples the test server, it accurately demonstrates security vulnerabilities, but it doesn't affect customers.

Remember this

A penetration test is an active test that can assess deployed security controls and determine the impact of a threat. It starts with a vulnerability scan and then tries to exploit vulnerabilities by actually attacking or simulating an attack.

White, Gray, and Black Box Testing

It's common to identify testing based on the level of knowledge the testers have prior to starting the test. These testers could be internal employees, or external security professionals working for a third-party organization to perform the test. The three types of testing are:

- **Black box testing.** Testers have zero knowledge of the environment prior to the test. Instead, they approach the test with the same knowledge as an attacker. When testing applications, black box testers wouldn't have any prior experience with the application. When testing networks, they aren't provided any information on the network before the test. This includes a lack of documentation and a lack of experience. Black box testers often use fuzzing to check for application vulnerabilities.
- **White box testing.** Testers have full knowledge of the environment. For example, they would have access to product documentation, source code, and possibly even logon details.
- **Gray box testing.** Testers have some knowledge of the environment but do not have access to all documentation or data.

Remember this

Black box testers have zero prior knowledge of the system prior to a penetration test. White box testers have full knowledge, and gray box testers have some knowledge. Black box testers often use fuzzing.

You may also come across the terms *black hat*, *white hat*, and *gray hat*. These aren't referring to testers but instead to different types of attackers. They are reminiscent of the Wild West, where you could easily identify the good guys and the bad guys by the color of their hat. Black hat identifies a malicious attacker performing criminal activities. White hat identifies a security professional working within the law. Gray hat identifies individuals who may have good intentions but their activities may cross ethical lines. For example, an activist, sometimes called a hacktivist, may use attack methods to further a cause, but not for personal gain.

Hackers and *crackers* are terms you may also come across. Originally, a hacker indicated someone proficient with computers who wanted to share

knowledge with others. They weren't malicious. In contrast, a cracker was a proficient hacker who used the knowledge for malicious purposes. However, English is a living language that continues to evolve and the media consistently uses the term hacker to identify malicious attackers. For clarity, this book uses the term *attacker* to identify an individual attacking a system for malicious purposes.

Obtaining Consent

It's important to obtain consent of the system owner before starting a penetration test. In most cases, this consent is in writing. If it isn't in writing, many security professionals won't perform the test. A penetration test without consent is an attack. An organization may perceive a well-meaning administrator doing an unauthorized penetration test as a black hat or gray hat attacker, and the administrator may soon be out of a job.

Many organizations use a written rules-of-engagement document when hiring outside security professionals to perform the test. The rules-of-engagement document identifies the boundaries of the penetration test. If testing does result in an outage even though the testers followed the rules of engagement, repercussions are less likely.

Intrusive Versus Nonintrusive Testing

Scans can be either intrusive or nonintrusive. An intrusive scan attempts to exploit vulnerabilities. In contrast, a nonintrusive scan attempts to determine if a vulnerability exists, but it does not try to exploit the vulnerability. You can also think of this as invasive and noninvasive, respectively.

Vulnerability scans are nonintrusive and less invasive than penetration tests. Penetration tests are intrusive and more invasive than vulnerability scans.

Passive Versus Active Tools

In the context of tools used to discover security threats and vulnerabilities, it's important to understand the difference between passive tools and active tools. A passive tool tests systems in a nonintrusive manner and has little possibility of compromising a system. An active tool uses intrusive and invasive methods and can potentially affect the operations of a system.

Throughout the “Vulnerability Assessment” section, it was stressed that vulnerability scanning is passive, whereas penetration testing is active. In this context, passive doesn't mean that a vulnerability scanner isn't doing anything. It certainly is probing systems to identify vulnerabilities and other problems. However, it does not take any action to exploit these vulnerabilities.

When preparing for any exam, including the CompTIA Security+ exam, it's worthwhile to look at the objectives. These objectives specifically use the word *passively*, and passive verbs, in the context of vulnerability scanning. They also use the word *actively*, and active verbs, in the context of penetration testing.

That doesn't mean that you can feel free to run a vulnerability scanner on any network because it is passive. If your actions are discovered, you can easily be identified as an attacker and face legal action.

Remember this

A vulnerability scanner is passive and nonintrusive and has little impact on a system during a test. In contrast, a penetration test is active and intrusive, and can potentially compromise a system. A pentest is more invasive than a vulnerability scan.

Continuous Monitoring

It's important to realize that there is never a time that security professionals can say, "Now that we've implemented this security measure, we can sit back knowing that we're safe." In other words, security is never finished. Instead, security professionals must continuously monitor their environment for emerging threats and new vulnerabilities.

Continuous security monitoring includes monitoring all relevant security controls, with the goal of ensuring that they help an organization maintain a strong security posture. There are many methods of monitoring, including performing periodic threat assessments, vulnerability assessments, and risk assessments. Many organizations perform routine vulnerability scans, such as once a week, and infrequent penetration tests. Additionally, organizations perform routine audits and reviews such as user rights and permissions reviews, which are discussed in the next section.

Identifying Security Tools

Several tools are available for use by security professionals and attackers alike. Vulnerability scanners were discussed at length earlier in this chapter, including their use as ping scanners and port scanners. However, other tools are available. This section discusses tools such as protocol analyzers, routine audits, and logs.

Sniffing with a Protocol Analyzer

A protocol analyzer can capture and analyze packets on a network. The process of using a protocol analyzer is sometimes referred to as sniffing or using a sniffer. Both administrators and attackers can use a protocol analyzer to view IP headers and examine packets. For example, administrators can use a protocol analyzer to troubleshoot communication issues between network systems, or identify potential attacks using manipulated or fragmented packets.

Attackers can use a protocol analyzer to capture data sent across a network in cleartext. One of the ways they do so is by connecting an unauthorized switch within a network to capture traffic and forward it to a system running a protocol analyzer. If cabling isn't protected, they might be able to simply connect a switch above a drop-down ceiling that wouldn't be detected easily.

Wireshark is a free protocol analyzer that you can download from here: <http://www.wireshark.org/>. Figure 8.1 shows Wireshark after it captured packets transmitted over the network. It includes about 150 packets and has packet 121 selected in the top pane. The top pane shows the source and destination IP addresses and the Server Message Block (SMB) protocol. Many networks use SMB to send files over the network, and this packet includes the contents of that file. The middle pane shows details from this packet with the Internet Protocol Version 4 header information partially expanded. The bottom pane shows the entire contents of the packet displayed in hexadecimal and ASCII characters.

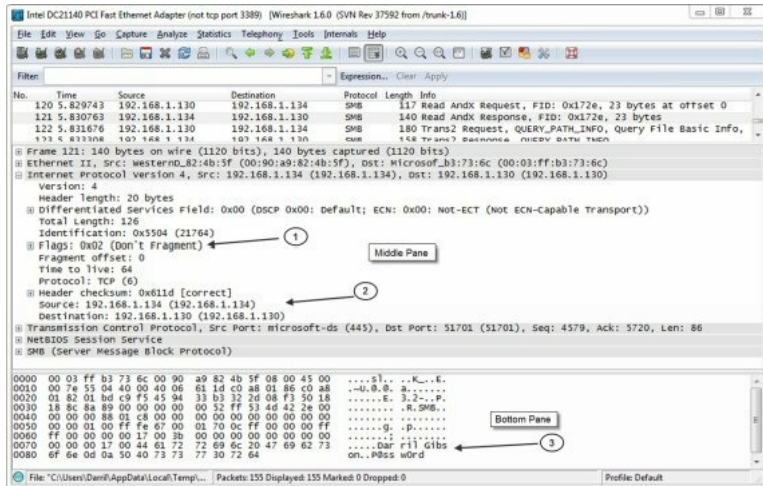


Figure 8.1: Wireshark capture

Although it can be tedious to analyze a packet capture, there is a lot of information in it for anyone willing to take the time to do so. Occasionally, attackers manipulate flags (arrow 1) within the headers for different types of attacks, and the protocol analyzer allows you to verify header manipulation attacks. You can also see the source and destination IP addresses (arrow 2) within the IP header field. You can expand the Ethernet II section to show the media access control (MAC) addresses of the source and destination computers.

Notice that you can view the username (Darril) and password (P@ssw0rd) in the bottom pane (arrow 3) because SMB sends it in cleartext. However, if an application encrypted the data before sending it across the network, it would not be readable.

Although this packet capture only includes about 150 packets, a packet capture can easily include thousands of packets. Wireshark includes filters that administrators use to focus on specific types of traffic. These filters also allow them to quantify the traffic. For example, they can determine the percentage of SMTP traffic or HTTP traffic on the network.

In addition to seeing a capture using the Wireshark graphical interface, you can also view them as text files (and that's how they'll typically be referenced on any exam). The information in the text file is usually limited using filters, but normally includes the time, source information labeled as src, destination information labeled as dst, and sometimes protocol information. Here's an example:

22:33:44, src 192.168.5.55:3389, dst 192.168.7.17:8080, syn/ack

The time is shown in a 24-hour clock as 10:33 p.m. and 44 seconds. Notice the source and destination include an IP address and a port number. This reiterates the importance of knowing the ports listed in Table 3.1. It also shows you how you can identify the source of traffic. For example, if an attacker is manipulating or fragmenting packets as part of an attack, you can use the src IP address to identify the potential source of the attack.

It's worth noting that the source IP address doesn't always identify the actual attacker. For example, attackers often take control of other computers and launch attacks from them without the knowledge of the owner. Similarly, Port Address Translation (PAT) translates public and private IP addresses. If the traffic goes through a device using PAT, the protocol analyzer only captures the translated IP address, not the original IP address.

When using a protocol analyzer, you need to configure the network interface card (NIC) on the system to use promiscuous mode. Normally, a NIC uses non-promiscuous mode and only processes packets addressed directly to its IP address. However when you put it in promiscuous mode, it processes all packets regardless of the IP address. This allows the protocol analyzer to capture all packets that reach the NIC.

Remember this

Administrators use a protocol analyzer (or sniffer) to capture, display, and analyze packets sent over a

network. It is useful when troubleshooting communications problems between systems. It is also useful to detect attacks that manipulate or fragment packets. A capture shows information such as the type of traffic (protocol), flags, source and destination IP addresses, and source and destination MAC addresses. The NIC must be configured to use promiscuous mode to capture all traffic.

Performing Routine Audits

Many organizations perform routine audits to help identify risks. An audit provides an independent and objective examination of processes and procedures. It can help an organization determine its security posture and verify that the organization is following its policies. Internal personnel or external auditors can conduct audits.

Routine audits can verify any system processes or organizational policy. For example, a security policy might state employees should notify a security manager as soon as possible after identifying an incident. The audit verifies personnel are following this policy. Another audit might verify that security controls such as a recently installed intrusion detection system (IDS) continues to function as intended.

Similarly, a security policy may state that when an employee leaves the company, an administrator must disable the employee's account. This prevents the ex-employee or someone else from using the account. This type of account management policy can prevent attacks such as the one at Fannie Mae (discussed in Chapter 6) where the employee installed a logic bomb after learning he lost his job.

If the audit discovers that accounts are not disabled, the organization then takes steps to identify and correct the problem. Is a written policy in place? Do appropriate personnel know their responsibilities in relation to the policy? Are processes in place that allow personnel to meet their responsibilities?

In some situations, existing processes don't support written policies, and the audit helps identify the problem. For example, imagine a single administrator is tasked with disabling accounts, but no one informs the administrator of employee terminations until days later. Obviously, there's no way the administrator will disable the accounts immediately. To prevent this type of situation, many organizations coordinate exit interviews with security personnel who disable the account during the exit interview.

User Reviews

User access reviews and user rights and permissions reviews are both a type of audit. A basic security principle is the principle of least privilege, and these reviews help verify users have the rights and permissions they need, but no more. This includes ensuring users have the ability to access only the resources they need to perform their job.

A user access review includes a review of what users have accessed, and can detect accounts from ex-employees that are still enabled. A user rights and permissions review identifies the privileges (rights and permissions) granted to users, and compares these against what the users need. These reviews can detect two common problems: privilege creep and inactive accounts.

Privilege creep (or permission bloat) occurs when a user is granted more and more privileges due to changing job requirements, but unneeded privileges are never removed. For example, imagine Lisa is working in the Human Resources (HR) department, so she has access to HR data. Later, she transfers to the Sales department and administrators grant her access to sales data. However, no one removes her access to HR data even though she doesn't need it to perform her job in the Sales department.

Organizations commonly use a role-based access control model with group-based privileges, as described in Chapter 2. For example, Lisa's user account would be in required HR department security groups, granting her appropriate privileges for her job. When she transfers, administrators would add her to the Sales department groups, granting her appropriate privileges for her new job. An organization should also have account management controls in place to ensure that administrators remove her account from the HR department security groups.

Most organizations ensure that user rights and permission reviews are performed at least once a year, and some organizations perform them more often. The goal is to do them often enough to catch potential problems and prevent security incidents. However, unless they can be automated, they become an unnecessary burden if security administrators are required to do them too often, such as daily or even once a week.

Remember this

Routine audits help an organization ensure they are following their policies, such as the principle of least privilege and account management control best practices. A user rights and permissions review ensures that users have only the access they need and no more and can detect privilege creep issues. It also ensures that inactive accounts are either disabled or deleted.

Monitoring Events with Logs

Logs have the capability to record what happened, when it happened, where it happened, and who did it. One of the primary purposes of logging is to allow someone, such as an administrator or security professional, to identify exactly what happened and when.

With this in mind, it's tempting to set up logging to record every event and provide as much detail as possible—most logs support a verbose mode that will log additional details. However, a limiting factor is the amount of disk space available. Additionally, when logging is enabled, there is an implied responsibility to review the logs. The more you choose to log, the more you may have to review. The following sections cover some commonly used logs.

Operating System Event Logs

Operating systems have basic logs that record events. For example, Windows systems have several common logs that record what happened on a Windows computer system. All of these logs are viewable using the Windows Event Viewer. One of the primary logs in a Windows system is the Security log and it functions as a security log, an audit log, and an access log.

The Security log records auditable events, such as when a user logs on or off, or when a user accesses a resource. Some auditing is enabled by default in some systems, but administrators can add additional auditing. The Security log records audited events as successes or failures. Success indicates an audited event completed successfully, such as a user successfully logging on or successfully deleting a file. Failure indicates that a user tried to perform an action but failed, such as failing to log on or trying to delete a file but receiving a permission error instead. Some additional logs in a Windows system include:

- **Application.** The Application log records events recorded by applications or programs running on the system. Any application has the capability of recording errors in the Application log.
- **System.** The operating system uses the System log to record events related to the functioning of the operating system. This can include when it starts, when it shuts down, information on services starting and stopping, drivers loading or failing, or any other system component event deemed important by the system developers.

If a system is attacked, you may be able to learn details of the attack by reviewing the operating system logs. Depending on the type of attack, any of the operating system logs may be useful.

Firewall and Router Access Logs

You can typically manipulate firewalls and routers to log specific information, such as logging all traffic that the device passes, all traffic that the device blocks, or both. These logs are useful when troubleshooting connectivity issues and when identifying potential intrusions or attacks.

Firewall and router logs include information on where the packet came from (the source) and where it is going (the destination). This includes IP addresses, MAC addresses, and ports.

Other Logs

In addition to the basic operating system logs and firewall and router access logs, administrators use other logs when maintaining systems and networks. These include:

- **Antivirus logs.** Antivirus logs log all antivirus activity, including when scans were run and if any malware was detected. These logs also identify if malware was removed or quarantined.
- **Application logs.** Many server applications include logging capabilities within the application. For example, database applications such as Microsoft SQL Server or Oracle Database include logs to record performance and user activity.
- **Performance logs.** Performance logs can monitor system performance and give an alert when preset performance thresholds are exceeded.

Reviewing Logs

Logs provide the ability to review activity, but ironically, this is often the most overlooked step in the auditing process. Often, administrators only dig into the logs when a symptom appears. Unfortunately, symptoms often don't appear until a problem has snowballed out of control.

Many third-party programs are available that can automate the review of logs for large organizations. For example, NetIQ has a full suite of applications that monitor multiple computers and servers in a network. When an event occurs, NetIQ examines the event to determine if it is an event of interest. If so, it triggers a programmed response, such as sending an email to a group of administrators.

Another benefit of a third-party program like this is that it provides centralized log management. If a system is attacked and compromised, the logs stored on the log server are retained. As a reminder, attackers often try to erase or modify logs after the attack. Centralized log management reduces the success of these attempts.

Chapter 8 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Identifying Risk

- A risk is the likelihood that a threat will exploit a vulnerability. A threat is a potential danger that can compromise confidentiality, integrity, or availability of data or a system. A vulnerability is a weakness.
- Risk management methods include risk avoidance, transference, acceptance, mitigation, and deterrence. You avoid a risk by not providing a service or not participating in a risky activity. Purchasing insurance, such as fire insurance, transfers the risk to another entity. Some controls such as security guards deter a risk.
- You cannot eliminate risk. Risk management attempts to reduce risk to a level that an organization is able to accept, and the remaining risk is known as residual risk. Senior management is responsible for managing risk and the losses associated from residual risk.
- Quantitative risk assessments use numbers, such as costs and asset values. The single loss expectancy (SLE) is the cost of any single loss. The annual rate of occurrence (ARO) indicates how many times the loss will occur annually. You can calculate the annual loss expectancy (ALE) as $SLE \times ARO$.
- Qualitative risk assessments use judgments to prioritize risks based on probability and impact. These judgments provide a subjective ranking.
- Risk assessment results are sensitive. Only executives and security professionals should be granted access to risk assessment reports.

Checking for Vulnerabilities

- A port scanner scans systems for open ports and attempts to discover what services and protocols are running.
- An advanced persistent threat typically attacks from another country and can launch sophisticated and targeted attacks.
- Vulnerability assessments determine the security posture of a system or network.
- Vulnerability scanners passively test security controls to identify vulnerabilities, a lack of security controls, and common misconfigurations. They are effective at discovering systems susceptible to an attack without exploiting the systems.
- A false positive from a vulnerability scan indicates the scan falsely detected a vulnerability, and the vulnerability doesn't actually exist.
- A penetration test is an active test that attempts to exploit discovered vulnerabilities. It starts with a vulnerability scan and then bypasses or actively tests security controls to exploit vulnerabilities.
- Significant differences between vulnerability scans and penetration tests are that vulnerability scans are passive and less invasive, while penetration tests are active and more invasive.
- A baseline review identifies changes from the original deployed configuration.
- Code reviews are a type of assessment where a peer programmer goes through code line-by-line looking for vulnerabilities, such as race conditions or susceptibility to buffer overflow attacks.
- Design reviews ensure that systems and software are developed properly, following standard security best practices.
- In black box testing, testers perform a penetration test with zero prior knowledge of the environment. White box testing indicates that the testers have full knowledge of the environment, including documentation and source code for tested applications. Gray box testing indicates some knowledge of the environment.
- Black hat indicates a malicious attacker, whereas white hat identifies a security professional working within the law.
- Penetration testers should gain consent prior to starting a penetration test. A rules-of-engagement document identifies the boundaries of the test.
- Continuous security monitoring helps an organization maintain its security posture, by verifying that security controls continue to function as intended.

Identifying Security Tools

- Protocol analyzers (sniffers) can capture and analyze data sent over a network. Attackers use protocol analyzers to capture cleartext data sent across a network.
- Administrators use protocol analyzers for troubleshooting communication issues by inspecting protocol headers to detect manipulated or fragmented packets.
- Captured packets show the type of traffic (protocol), source and destination IP addresses, source and destination MAC addresses, and flags.
- Routine audits help an organization verify they are following their own policies, such as the principle of least privilege and account management control best practices.
- A user rights and permissions review ensures that users have only the access they need and no more. It also verifies that inactive accounts are either disabled or deleted.

- Security logs track logon and logoff activity on systems. System logs identify when services start and stop.
- Firewall and router logs identify the source and destination of traffic.
- Centralized log management protects logs when systems are attacked or compromised.

Chapter 8 Practice Questions

1. Which of the following is most closely associated with residual risk?
 - A. Risk acceptance
 - B. Risk avoidance
 - C. Risk deterrence
 - D. Risk mitigation
 - E. Risk transference
2. You need to calculate the ALE for a server. The value of the server is \$3,000, but it has crashed 10 times in the past year. Each time it crashed, it resulted in a 10 percent loss. What is the ALE?
 - A. \$300
 - B. \$500
 - C. \$3,000
 - D. \$30,000
3. You need to calculate the expected loss of an incident. Which of the following value combinations would you MOST likely use?
 - A. ALE and ARO
 - B. ALE and SLE
 - C. SLE and ARO
 - D. ARO and ROI
4. You want to identify all of the services running on a server. Which of the following tools is the BEST choice to meet this goal?
 - A. Penetration test
 - B. Protocol analyzer
 - C. Sniffer
 - D. Port scanner
5. You recently completed a vulnerability scan on your network. It reported that several servers are missing key operating system patches. However, after checking the servers, you've verified the servers have these patches installed. Which of the following BEST describes this?
 - A. False negative
 - B. Misconfiguration on servers
 - C. False positive
 - D. Servers not hardened
6. You suspect that a database server used by a web application does not have current patches. Which of the following is the BEST action to take to verify the server has up-to-date patches?
 - A. Vulnerability scan
 - B. Port scan
 - C. Protocol analyzer
 - D. Host enumeration
7. You need to perform tests on your network to identify missing security controls. However, you want to have the least impact on systems that users are accessing. Which of the following tools is the best to meet this need?
 - A. Code review
 - B. Vulnerability scan
 - C. Ping sweep
 - D. Penetration test
8. Lisa needs to identify if a risk exists on a web application and if attackers can potentially bypass security controls. However, she should not actively

test the application. Which of the following is the BEST choice?

- A. Perform a penetration test.
- B. Perform a port scan.
- C. Perform a vulnerability scan.
- D. Perform traffic analysis with a sniffer.

9. A recent vulnerability scan reported that a web application server is missing some patches. However, after inspecting the server, you realize that the patches are for a protocol that administrators removed from the server. Which of the following is the BEST explanation for this disparity?

- A. False negative
- B. False positive
- C. Lack of patch management tools
- D. The patch isn't applied

10. Your organization develops web application software, which it sells to other companies for commercial use. Your organization wants to ensure that the software isn't susceptible to common vulnerabilities, such as buffer overflow attacks and race conditions. What should the organization implement to ensure software meets this standard?

- A. Input validation
- B. Change management
- C. Code review
- D. Regression testing

11. An organization has a legacy server within the DMZ. It is running older software that is not compatible with current patches, so it remains unpatched. Management accepts the risk on this system, but wants to know if attackers can access the internal network if they successfully compromise this server. Which of the following is the MOST appropriate test?

- A. Vulnerability scan
- B. Port scan
- C. Code review
- D. Pentest

12. Testers do not have access to product documentation or any experience with an application. What type of test will they MOST likely perform?

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

13. Your organization has hired a group of external testers to perform a black box penetration test. One of the testers asks you to provide information about your internal network. What should you provide?

- A. A list of IP ranges and the types of security devices operational on the network
- B. Network diagrams but without internal IP addresses
- C. Some network diagrams and some IP addresses, but not all
- D. Nothing

14. A network administrator is troubleshooting a communication problem between a web server and a database server. Which of the following tools would MOST likely be useful in this scenario?

- A. Protocol analyzer
- B. Port scanner
- C. Switch
- D. URL filter

15. A network administrator needs to identify the type of traffic and packet flags used in traffic sent from a specific IP address. Which of the following is the BEST tool to meet this need?

- A. UTM security appliance

- B. Router logs
- C. Protocol analyzer
- D. Vulnerability scan

16. While analyzing a packet capture log, you notice the following entry:

16:12:50, src 10.80.1.5:3389, dst 192.168.1.100:8080, syn/ack

Of the following choices, what is the BEST explanation of this entry?

- A. An HTTP connection attempt
- B. An RDP connection attempt
- C. An FTP connection attempt
- D. A buffer overflow attack

17. Security administrators have recently implemented several security controls to enhance the network’s security posture. Management wants to ensure that these controls continue to function as intended. Which of the following tools is the BEST choice to meet this goal?

- A. Routine audit
- B. Change management
- C. Design review
- D. Black box test

18. Your organization recently hired an outside security auditor to review internal processes. The auditor identified several employees who had permissions for previously held jobs within the company. What should the organization implement to prevent this in the future?

- A. Design reviews
- B. Code reviews
- C. Baseline review
- D. User rights and permissions reviews

19. Your organization’s security policy states that administrators should follow the principle of least privilege. Which of the following tools can ensure that administrators are following the policy?

- A. User rights and permissions review
- B. Risk assessment
- C. Vulnerability assessment
- D. Threat assessment

20. Your organization wants to ensure that security controls continue to function, helping to maintain an appropriate security posture. Which of the following is the BEST choice to meet this goal?

- A. Auditing logs
- B. Routine audits
- C. Continuous security monitoring
- D. Vulnerability scans

Chapter 8 Practice Question Answers

1. **A.** Residual risk is the risk that an organization accepts after implementing controls to reduce risk. An organization can avoid a risk by not providing a service or not participating in a risky activity. Risk deterrence attempts to discourage attacks with preventive controls such as a security guard. Risk mitigation reduces risks through internal controls. Purchasing insurance is a common method of risk transference.
2. **C.** The annual loss expectancy (ALE) is \$3,000. It is calculated as single loss expectancy (SLE) \times annual rate of occurrence (ARO). The SLE is 10 percent of \$3,000 (\$300) and the ARO is 10. $10 \times \$300$ is \$3,000.
3. **A.** The expected loss is the single loss expectancy (SLE) and you can calculate it with the annual loss expectancy (ALE) and annual rate of occurrence (ARO), as ALE / ARO . The SLE is what you are trying to determine, so you don't have that value. The return on investment (ROI) will not help in identifying the SLE.
4. **D.** A port scanner identifies open ports on a system and is commonly used to determine what services are running on the system. A penetration test attempts to exploit a vulnerability. A protocol analyzer (also called a sniffer) could analyze traffic and discover protocols in use, but this would be much more difficult than using a port scanner.
5. **C.** In this scenario, the vulnerability scanner reported a false positive indicating that the servers had a vulnerability, but in reality, the servers did not have the vulnerability. A false negative occurs if a vulnerability scanner does not report a known vulnerability. There isn't any indication that the servers are misconfigured and they are not hardened.
6. **A.** A vulnerability scan determines if the system has current patches and is the best choice of those given. A port scan identifies open ports. A protocol analyzer (sniffer) captures traffic for analysis. Host enumeration identifies hosts on a network based on their IP addresses.
7. **B.** A vulnerability scanner is passive and has the least impact on systems, but it can detect systems that are lacking specific security controls. A code review is effective for identifying vulnerabilities in software. However, it doesn't identify missing security controls elsewhere. A ping sweep can identify hosts on a network based on their IP addresses. A penetration test does not have the least impact on systems.
8. **C.** A vulnerability scan identifies vulnerabilities that attackers can potentially exploit, and vulnerability scanners perform passive testing. A penetration test actively tests the application and can potentially compromise the system. A port scan only identifies open ports. A sniffer can capture traffic for analysis, but it doesn't check for security controls.
9. **B.** A false positive on a vulnerability scan indicates that a vulnerability is positively detected, but the vulnerability doesn't actually exist. A false negative indicates that the vulnerability scan did not detect a vulnerability that does exist on a system. False positives can occur even if an organization has a strong patch management process in place. Although it's true that the patch isn't applied, it's also true that the patch cannot be applied because it is for a protocol that administrators removed.
10. **C.** A code review goes line-by-line through the software code looking for vulnerabilities, such as buffer overflows and race conditions. Input validation helps prevent buffer overflows but not race conditions. Change management controls help prevent unintended outages from unauthorized changes. Regression testing is a type of testing used to ensure that new patches do not cause errors.
11. **D.** A pentest (or penetration test) attempts to compromise the server and then attempts to access the internal network. A vulnerability scan is passive. It does not attempt to compromise a system, so it cannot verify if an attacker can access the internal network. A port scan only identifies open ports. A code review is useful for newly developed software, but there isn't any indication that the original code is available for the legacy server.
12. **C.** A black box tester does not have access to product documentation or experience with an application. White box testers have full knowledge and gray box testers have some knowledge. Black hat refers to a malicious attacker.
13. **D.** Black box testers should not have access to any information before starting the test, so technicians and administrators should not provide any information if asked. It's appropriate to give white box testers all the information on the network, and give gray box testers some information on the internal network.
14. **A.** A protocol analyzer (or sniffer) is useful for capturing traffic between systems for analysis and is the best choice for this scenario. A port scanner identifies open ports in single systems, so it wouldn't be helpful here. Traffic between the systems likely goes through the switch and you can monitor traffic going through the switch with a protocol analyzer, but by itself, the switch wouldn't help troubleshoot a communication problem. A URL filter filters outgoing web traffic.
15. **C.** A protocol analyzer (or sniffer) can capture traffic sent over a network and identify the type of traffic, the source of the traffic, and protocol flags used within individual packets. A unified threat management (UTM) security appliance combines multiple security solutions into a single solution but doesn't typically capture traffic. Router logs identify the type of traffic going through it, but do not include packet flag data. A vulnerability scan identifies vulnerabilities on a network.
16. **B.** This log entry indicates that a source (src) system with an IP of 10.80.1.5 sent a connection attempt using port 3389, which is the Remote Desktop Protocol (RDP) port, at time 4:12:50 p.m. The destination (dst) was sent to IP 192.168.1.100 using a common proxy server listening port of 8080. Hypertext Transfer Protocol (HTTP) uses port 80, not port 3389. File Transfer Protocol (FTP) uses ports 20 and 21, not port 3389. A buffer overflow attack sends unexpected data, but this entry indicates that it is a SYN/ACK (synchronize/acknowledge) packet establishing a connection.

17. **A.** A routine audit can verify controls are continuing to operate as intended. Change management controls can help ensure that systems don't suffer from unintended outages after a change, and although change management helps ensure the controls aren't modified, it doesn't necessarily ensure the controls continue to operate as intended. A design review would be done before the controls are deployed. A black box test is a type of penetration test where the testers don't have any knowledge of the system, so it wouldn't be able to identify if the controls are functioning as intended.
18. **D.** A user rights and permissions review detects permission bloat situations such as this. Account management controls also help ensure these situations don't occur. A design review helps ensure that systems and software are developed properly. A code review is a line-by-line review of code by peer programmers. A baseline review compares current configurations against baseline settings.
19. **A.** A user rights and permissions review verifies users have the permissions they need for their job, and no more, which verifies the principle of least privilege is being followed. Risk, vulnerability, and threat assessments assess current risks, and they might verify the principle of least privilege is being followed, but they do much more.
20. **C.** Continuous security monitoring helps an organization maintain its security posture, by verifying that security controls continue to function as intended. Auditing logs, performing routine audits, and performing vulnerability scans are all part of a continuous monitoring plan. However, individually, they do not verify all security controls are operating properly.

Chapter 9

Preparing for Business Continuity

CompTIA Security+ objectives covered in this chapter:

1.1 Implement security configuration parameters on network devices and other technologies.

- Load Balancers

2.1 Explain the importance of risk related concepts.

- Recovery time objective and recovery point objective

2.3 Given a scenario, implement appropriate risk mitigation strategies.

- Enforce policies and procedures to prevent data loss or theft

2.7 Compare and contrast physical security and environmental controls.

- Environmental controls (HVAC, Fire suppression, EMI shielding, Hot and cold aisles, Environmental monitoring, Temperature and humidity controls)
- Physical security, Protected distribution (cabling)

2.8 Summarize risk management best practices.

- Business continuity concepts (Business impact analysis, Identification of critical systems and components, Removing single points of failure, Business continuity planning and testing, Continuity of operations, Disaster recovery, IT contingency planning, Succession planning, High availability, Redundancy, Tabletop exercises)
- Fault tolerance (Hardware, RAID, Clustering, Load balancing, Servers)
- Disaster recovery concepts (Backup plans/policies, Backup execution/frequency, Cold site, Hot site, Warm site)

2.9 Given a scenario, select the appropriate control to meet the goals of security.

- Availability (Redundancy, Fault tolerance)
- Safety (Escape plans, Drills, Escape routes, Testing controls)

**

Although you can't prevent some disasters, such as hurricanes or floods, you can prevent catastrophic failures by taking preventive steps. You can identify single points of failure, implement redundancy solutions, and perform backups. Many organizations use formal business continuity and disaster recovery plans to prepare for potential disasters. This chapter covers these concepts and some key environmental controls.

Adding Redundancy

One of the constants with computers, subsystems, and networks is that they will fail. It's one of the few things you can count on. It's not a matter of *if* they will fail, but *when*. However, by adding redundancy into your systems and networks, you can increase the reliability of your systems even when they fail. By increasing reliability, you increase one of the core security goals: availability.

Redundancy adds duplication to critical system components and networks and provides fault tolerance. If a critical component has a fault, the duplication provided by the redundancy allows the service to continue as if a fault never occurred. In other words, a system with fault tolerance can suffer a fault, but it can tolerate it and continue to operate. Organizations often add redundancies to eliminate single points of failure.

You can add redundancies at multiple levels:

- Disk redundancies using RAID
- Server redundancies by adding failover clusters
- Power redundancies by adding generators or a UPS
- Site redundancies by adding hot, cold, or warm sites

Single Point of Failure

A *single point of failure* is a component within a system that can cause the entire system to fail if the component fails. When designing redundancies, an organization will examine different components to determine if they are a single point of failure. If so, they take steps to provide a redundancy or fault-tolerance capability. The goal is to increase reliability and availability of the systems.

Some examples of single points of failure include:

- **Disk.** If a server uses a single drive, the system will crash if the single drive fails. Redundant array of inexpensive disks (RAID) provides fault tolerance for hard drives and is a relatively inexpensive method of adding fault tolerance to a system.
- **Server.** If a server provides a critical service and its failure halts the service, it is a single point of failure. Failover clusters (discussed later in this chapter) provide fault tolerance for critical servers.
- **Power.** If an organization only has one source of power for critical systems, the power is a single point of failure. However, elements such as uninterruptible power supplies (UPSs) and power generators provide fault tolerance for power outages.

Although information technology (IT) personnel recognize the risks with single points of failure, they often overlook them until a disaster occurs.

However, tools such as business continuity plans (covered later in this chapter) help an organization identify critical services and address single points of failure.

Remember this

A single point of failure is any component whose failure results in the failure of an entire system. Elements such as RAID, failover clustering, UPSs, and generators remove many single points of failure. RAID is an inexpensive method used to add fault tolerance and increase availability.

Disk Redundancies

Any system has four primary resources: processor, memory, disk, and the network interface. Of these, the disk is the slowest and most susceptible to failure. Because of this, administrators often upgrade disk subsystems to improve their performance and redundancy.

Redundant array of inexpensive disks (RAID) subsystems provide fault tolerance for disks and increase the system availability. Even if a disk fails, most RAID subsystems can tolerate the failure and the system will continue to operate. RAID systems are becoming much more affordable as the price of drives steadily falls and disk capacity steadily increases.

RAID-0

RAID-0 (striping) is somewhat of a misnomer because it doesn't provide any redundancy or fault tolerance. It includes two or more physical disks. Files stored on a RAID-0 array are spread across each of the disks.

The benefit of a RAID-0 is increased read and write performance. Because a file is spread across multiple physical disks, the different parts of the file can be read from or written to each of the disks at the same time. If you have three 500 GB drives used in a RAID-0, you have 1500 GB (1.5 TB) of storage space.

RAID-1

RAID-1 (mirroring) uses two disks. Data written to one disk is also written to the other disk. If one of the disks fails, the other disk still has all the data, so the system can continue to operate without any data loss. With this in mind, if you mirror all the drives in a system, you can actually lose half of the drives and continue to operate.

You can add an additional disk controller to a RAID-1 configuration to remove the disk controller as a single point of failure. In other words, each of the disks also has its own disk controller. Adding a second disk controller to a mirror is called disk duplexing.

If you have two 500 GB drives used in a RAID-1, you have 500 GB of storage space. The other 500 GB of storage space is dedicated to the fault-tolerant, mirrored volume.

RAID-2, RAID 3, and RAID-4 are rarely used.

RAID-5 and RAID-6

A RAID-5 is three or more disks that are striped together similar to RAID-0. However, the equivalent of one drive includes parity information. This parity information is striped across each of the drives in a RAID-5 and is used for fault tolerance. If one of the drives fails, the system can read the information on the remaining drives and determine what the actual data should be. If two of the drives fail in a RAID-5, the data is lost.

RAID-6 is an extension of RAID-5, and it includes an additional parity block. A huge benefit is that the RAID-6 disk subsystem will continue to operate even if two disk drives fail. RAID-6 requires a minimum of four disks.

Remember this

RAID subsystems, such as RAID-1, RAID-5, and RAID-6, provide fault tolerance and increased data availability. RAID-5 can survive the failure of one disk. RAID-6 can survive the failure of two disks.

RAID-10

A RAID-10 configuration combines the features of mirroring (RAID-1) and striping (RAID-0). RAID-10 is sometimes called RAID 1+0. A variation is RAID-01 or RAID 0+1 that also combines the features of mirroring and striping but implements the drives a little differently.

Software Versus Hardware RAID

Hardware RAID configurations are significantly better than software RAID. In hardware RAID, dedicated hardware manages the disks in the RAID, removing the load from the operating system. In contrast, the operating system manages the disks in the RAID array in software RAID. Hardware RAID systems provide better overall performance and often include extra features.

For example, a hardware RAID may include six physical disks using four in an active RAID-6 configuration and two as online spares. If one of the active disks in the RAID-6 fails, the RAID will continue to operate because a RAID-6 can tolerate the failure.

However, a hardware RAID can logically take the failed disk out of the configuration, add one of the online spares into the configuration, and rebuild the array. All of this happens without any administrator intervention. Hardware RAID systems are often hot swappable, allowing administrators to swap out the failed drive without powering the system down.

Server Redundancy

Some services require a high level of availability and it's possible to achieve 99.999 percent uptime, commonly called five nines. It equates to less than 6 minutes of downtime a year: $60 \text{ minutes} \times 24 \text{ hours} \times 365 \text{ days} \times .00001 = 5.256 \text{ minutes}$. Failover clusters are a key component used to achieve five nines.

Although five nines is achievable, it's expensive. However, if the potential cost of an outage is high, the high cost of the redundant technologies is justified. For example, some web sites generate a significant amount of revenue, and every minute a web site is unavailable represents lost money. High-capacity failover clusters ensure the service is always available even if a server fails.

Failover Clusters for High Availability

The primary purpose of a failover cluster is to provide high availability for a service offered by a server. Failover clusters use two or more servers in a cluster configuration, and the servers are referred to as nodes. At least one server or node is active and at least one is inactive. If an active node fails, the inactive node can take over the load without interruption to clients.

Consider Figure 9.1, which shows a two-node failover cluster. Both nodes are individual servers, and they both have access to external data storage used by the active server. Additionally, the two nodes have a monitoring connection to each other used to check the health or heartbeat of each other.




Figure 9.1: Failover cluster

Imagine that Node 1 is the active node. When any of the clients connect, the cluster software (installed on both nodes) ensures that the clients connect to the active node. If Node 1 fails, Node 2 senses the failure through the heartbeat connection and configures itself as the active node. Because both nodes have access to the shared storage, there is no loss of data for the client. Clients may notice a momentary hiccup or pause, but the service continues.

You might notice that the shared storage in Figure 9.1 represents a single point of failure. It's not uncommon for this to be a robust hardware RAID-6. This ensures that even if two hard drives in the shared storage fails, the service will continue. Additionally, if both nodes are plugged into the same power grid, the power represents a single point of failure. They can each be protected with a separate uninterruptible power supply (UPS), and use a separate power grid.

Cluster configurations can include many more nodes than just two. However, nodes need to have close to identical hardware and are often quite expensive, but if a company truly needs to achieve 99.999 percent uptime, it's worth the expense.

Load Balancers for High Availability

A load balancer can optimize and distribute data loads across multiple computers or multiple networks. For example, if an organization hosts a popular web site, it can use multiple servers hosting the same web site in a web farm. Load-balancing software distributes traffic equally among all the servers in the web farm.

The term *load balancer* makes it sound like it's a piece of hardware, but a load balancer can be hardware or software. A hardware-based load balancer accepts traffic and directs it to servers based on factors such as processor utilization and the number of current connections to the server. A software-based load balancer uses software running on each of the servers in the load-balanced cluster to balance the load.

Load balancing primarily provides scalability, but it also contributes to high availability. Scalability refers to the ability of a service to serve more clients without any decrease in performance. Availability ensures that systems are up and operational when needed. By spreading the load among multiple systems, it ensures that individual systems are not overloaded, increasing overall availability.

Consider a web server that can serve 100 clients per minute, but if more than 100 clients connect at a time, performance degrades. You need to either scale up or scale out to serve more clients. You scale the server up by adding additional resources, such as processors and memory, and you scale out by adding additional servers in a load balancer.

Figure 9.2 shows an example of a load balancer with multiple web servers. Each web server includes the same web application. Some load balancers simply send new clients to the servers in a round-robin fashion. The load balancer sends the first client to Server 1, the second client to Server 2, and so on. Other load balancers automatically detect the load on individual servers and send new clients to the least used server.




Figure 9.2: Load balancing

An added benefit of many load balancers is that they can detect when a server fails. If a server stops responding, the load-balancing software no longer sends clients to this server. This contributes to overall high availability for the load balancer.

When servers are load balanced, it's called a load-balanced cluster, but it is not the same as a failover cluster. A failover cluster provides high availability by ensuring another node can pick up the load for a failed node. A load-balanced cluster provides high availability by sharing the load among multiple servers. When systems must share the same data storage, a failover cluster is appropriate. However, when the systems don't need to share the same storage, a load-balancing solution is more appropriate, and less expensive. Also, it's relatively easy to add additional servers to a load-balancing solution.

Remember this

Failover clusters are one method of server redundancy and they provide high availability for servers. They can remove a server as a single point of failure. Load balancing increases the overall processing power of a service by sharing the load among multiple servers. Load balancers also ensure availability when a service has an increased number of requests.

Power Redundancies

Power is a critical utility to consider when reviewing redundancies. For mission-critical systems, you can use uninterruptible power supplies and generators to provide both fault tolerance and high availability.

UPS

An uninterruptible power supply (UPS) is a battery or bank of batteries used as a backup in case of primary power failure. The UPS plugs into the wall and receives power from a commercial power source. The commercial power keeps the batteries charged and electronics within the UPS system provide power to external systems.

If commercial power fails, the UPS continues to provide power without any interruption. One of the added benefits of a UPS system is that it can protect against power fluctuations. Even if commercial power has momentary fluctuations, external systems aren't affected because they receive their power directly from the UPS.

Common UPS systems provide power for 10 to 15 minutes after a power outage. They aren't meant to provide longer-term power. Instead, the goal is to provide power until one of the following events occurs:

- **The supported system has enough time to shut down.** For example, a 10-minute UPS may send a shutdown signal to a system after power has been lost for 5 minutes. The system now has 5 minutes to perform an orderly shutdown.
- **Generators have enough time to power up and stabilize.** Both UPSs and generators support critical systems. The UPS provides short-term power, and the generator provides long-term power.
- **Commercial power returns.** A UPS provides fault tolerance for short outages and momentary power fluctuations. When commercial power returns, it recharges the batteries to get them back to full potential.

Generators

Many organizations use generators for critical systems that need long-term power. It isn't feasible to keep the generators running all the time due to fuel costs. Instead, they are started when power fails. Because it takes time for a generator to rev up to full power and stabilize, it cannot provide AC power immediately. A UPS powers critical systems until the generators stabilize.

In some cases, a generator automatically turns on when it detects power is lost and automatically switches over to generator power after the generator stabilizes. In other cases, technicians power up the generators manually, and manually switch over to generator power when the generators stabilize.

Remember this

An uninterruptible power supply (UPS) provides fault tolerance for power and can protect against power fluctuations. A UPS provides short-term power. Generators provide long-term power in extended outages.

Protecting Data with Backups

Backups are copies of data created to ensure that if the original data is lost or corrupted, it can be restored. Maybe I should restate that. Backups are copies of data created to ensure that *when* the original data is lost or corrupted, it can be restored. The truth is, if you work with computers long enough, you will lose data. The difference between a major catastrophe and a minor inconvenience is the existence of a backup.

It's important to realize that redundancy and backups are not the same thing. Protecting data with a RAID-1 or RAID-6 does not negate the need for backups. If a fire destroys a server, it also destroys the data on the RAID. Without a backup, all of the data is gone. Forever.

A Backup Horror Story

A friend of mine was a consultant for small businesses and was once hired to help a small business owner recover some lost data. The owner had been growing his business for about five years and had just about everything related to his business (client lists, billing information, proposals, agreements, and more) on one system. This system crashed.

The consultant tried to restore information from the disk but couldn't restore any data. The business owner panicked, knowing he simply needed the information. If he couldn't get the data back, his business might fail.

Although it's expensive, it is possible to have a clean-room facility take a hard drive apart and read the data at the bit level to restore at least some of the data. At this point, the owner was willing to try anything, so he paid the high price and they sent the disk to a recovery facility. Unfortunately, the disk suffered a catastrophic failure, and they weren't able to retrieve any meaningful data even in the clean room.

My friend visited the owner to relay the bad news. He said that when he left, the owner had his head in his hands and was literally crying. The business he had built for five years was close to ruins without much chance for recovery.

The worst part of this story is that it's repeated over and over with many different people in many different environments. As an example of a larger-scale disaster, MegaPetCo was a small chain of pet stores enjoying success when a single failure bankrupted them in 2009. An administrator accidentally performed a bulk update that deleted all the data in their primary database. There were zero backups. Within a few months, they filed for bankruptcy, closed all their stores, and laid off several hundred employees.

Too many people don't recognize the importance of backups until they've lost their data. Unfortunately, by then, it's too late.

...

Comparing Backup Types

Backup utilities support several different types of backups. Even though third-party backup programs can be quite sophisticated in what they do and how they do it, you should have a solid understanding of the basics.

The most common media used for backups is tape. Tapes store more data and are cheaper than other media, though some organizations use hard disk drives for backups. However, the type of media doesn't affect the backup type.

The following backup types are the most common:

- **Full backup.** A full (or normal backup) backs up all the selected data.
- **Differential backup.** This backs up all the data that has changed or is different since the last full backup.
- **Incremental backup.** This backs up all the data that has changed since the last full or incremental backup.

Full Backups

A full backup backs up all data specified in the backup. For example, you could have several folders on the D: drive. If you specify these folders in the backup program, the backup program backs up all the data in these folders.

Although it's possible to do a full backup on a daily basis, it's rare to do so in most production requirements. This is because of two limiting factors:

- **Time.** A full backup can take several hours to complete and can interfere with operations. However, administrators don't always have unlimited time to do backups and other system maintenance. For example, if a system is online 24/7, administrators may need to limit the amount of time for full backups to early Sunday morning to minimize the impact on users.
- **Money.** Backups need to be stored on some type of media, such as tape or hard drives. Performing full backups every day requires more media, and the cost can be prohibitive.

Instead, organizations often combine full backups with differential or incremental backups. However, every backup strategy must start with a full backup.

Restoring a Full Backup

A full backup is the easiest and quickest to restore. You only need to restore the single full backup and you're done. If you store backups on tapes, you only need to restore a single tape. However, most organizations need to balance time and money and use either a full/differential or a full/incremental backup strategy.

Differential Backups

A differential backup strategy starts with a full backup. After the full backup, differential backups back up data that has changed or is different since the last full backup.

For example, a full/differential strategy could start with a full backup on Sunday night. On Monday night, a differential backup would back up all files that changed since the last full backup on Sunday. On Tuesday night, the differential backup would again back up all the files that changed since the last full backup. This repeats until Sunday, when another full backup starts the process again. As the week progresses, the differential backup steadily grows in size.

Restoring a Full/Differential Backup Set

Assume for a moment that each of the backups was stored on different tapes. If the system crashed on Wednesday morning, how many tapes would you need to recover the data?

The answer is two. You would first recover the full backup from Sunday. Because the differential backup on Tuesday night includes all the files that changed after the last full backup, you would restore that tape to restore all the changes up to Tuesday night.

Incremental Backups

An incremental backup strategy also starts with a full backup. After the full backup, incremental backups then back up data that has changed since the last backup. This includes either the last full backup, or the last incremental backup.

As an example, a full/incremental strategy could start with a full backup on Sunday night. On Monday night, an incremental backup would back up all the files that changed since the last full backup. On Tuesday night, the incremental backup would back up all the files that changed since the incremental backup on Monday night. Similarly, the Wednesday night backup would back up all files that changed since the last incremental backup on Tuesday night. This repeats until Sunday when another full backup starts the process again. As the week progresses, the incremental backups stay about

the same size.

Restoring a Full/Incremental Backup Set

Assume for a moment that each of the backups was stored on a different tape. If the system crashed on Thursday morning, how many tapes would you need to recover the data?

The answer is four. You would first need to recover the full backup from Sunday. Because the incremental backups would be backing up different data each day of the week, each of the incremental backups must be restored and in the chronological order.

Sometimes, people mistakenly think the last incremental backup would have all the relevant data. Although it might have some relevant data, it doesn't have everything.

As an example, imagine you worked on a single project file each day of the week, and the system crashed on Thursday morning. In this scenario, the last incremental backup would hold the most recent copy of this file. However, what if you compiled a report every Monday but didn't touch it again until the following Monday? Only the incremental backup from Monday would include the most recent copy. An incremental backup from Wednesday night or another day of the week wouldn't include the report.

Choosing Full/Incremental or Full/Differential

A logical question is, "Why are there so many choices for backups?" The answer is that different organizations have different needs.

For example, imagine two organizations perform daily backups to minimize losses. They each do a full backup on Sunday, but are now trying to determine if they should use a full/incremental or a full/differential strategy.

The first organization doesn't have much time to perform maintenance throughout the week. In this case, the backup administrator needs to minimize the amount of time required to complete backups during the week. An incremental backup only backs up the data that has changed since the last backup. In other words, it includes changes only from a single day. In contrast, a differential backup includes all the changes since the last full backup. Backing up the changes from a single day takes less time than backing up changes from multiple days, so a full/incremental backup is the best choice.

In the second organization, recovery of failed systems is more important. If a failure requires restoring data, they want to minimize the amount of time needed to restore the data. A full/differential is the best choice in this situation because it only requires the restoration of two backups, the full and the most recent differential backup. In contrast, a full/incremental can require the restoration of several different backups, depending on when the failure occurs.

Remember this

If you have unlimited time and money, the full backup alone provides the fastest recovery time. Full/incremental strategies reduce the amount of time needed to perform backups. Full/differential strategies reduce the amount of time needed to restore backups.

Testing Backups

I've heard many horror stories in which personnel are regularly performing backups thinking all is well. Ultimately, something happens and they need to restore some data. Unfortunately, they discover that none of the backups holds valid data. People have been going through the motions, but something in the process is flawed.

The only way to validate a backup is to perform a test restore. Performing a *test restore* is nothing more than restoring the data from a backup and verifying its integrity. If you want to verify that you can restore the entire backup, you perform a full restore of the backup. If you want to verify that you can restore individual files, you perform a test restore of individual files. It's common to restore data to a different location other than the original source location, but in such a way that you can validate the data.

As a simple example, an administrator can retrieve a random backup and attempt to restore it. There are two possible outcomes of this test, and both are good:

- **The test succeeds.** Excellent! You know that the backup process works. You don't necessarily know that every backup tape is valid, but at least you know that the process is sound and at least some of your backups work.
- **The test fails.** Excellent! You know there's a problem that you can fix before a crisis. If you discovered the problem after you actually lost data, it wouldn't help you restore the data.

An additional benefit of performing regular test restores is that it allows administrators to become familiar with the process. The first time they do a restore shouldn't be in the middle of a crisis with several high-level managers peering over their shoulders.

Protecting Backups

If data is important enough to be backed up, it's important enough to protect. Backup media should be protected at the same level as the data that it holds. In other words, if proprietary data enjoys the highest level of protection within an organization, then backups of this data should also have the highest level of protection.

Protecting backups includes:

- **Storage.** This includes using clear labeling to identify the data and physical security protection to prevent others from easily accessing it while it's stored.
- **Transfer.** Data should be protected any time it is transferred from one location to another. This is especially true when transferring a copy of the backup to a separate geographical location.
- **Destruction.** When the backups are no longer needed, they should be destroyed. This can be accomplished by degaussing the media, shredding or burning the media, or scrubbing the media by repeatedly writing varying patterns of 1s and 0s onto the media.

Remember this

Test restores are the best way to test the integrity of a company's backup data. Backup media should be protected with the same level of protection as the data on the backup.

Backup Policies and Plans

Organizations typically create a backup policy to answer critical questions related to backups. Once the backup policy is created, administrators then implement backup plans to meet the needs addressed in the backup policy.

Unfortunately, many organizations operate without a backup policy, making it difficult for administrators to create appropriate backup plans. If they don't do any backups, management will blame them when data is lost. More than a few administrators have asked managers, "How much data are you willing to lose?" and heard "None!" as the response. However, if they try to back up everything and keep it forever, management will blame them for spending too much money. The ideal solution is somewhere in the middle, but where? How many backups and how much money is appropriate?

The backup policy is a written document and will often include the following details:

- **Identifies data to backup.** This identifies data important enough to back up. When management doesn't identify this in a backup policy, administrators and technicians have to make individual decisions on what data they consider important. Their decisions might not match management's value of the data.
- **Requires off-site backups.** A copy of a backup should be stored in a separate geographical location. This protects against a disaster such as a fire or flood. Even if a disaster destroys the site, the organization will still have another copy of the critical data.
- **Requires labeling media.** Media labels identify the data and the date of the backup. When data needs to be restored, the administrator should be able to quickly identify the backup that holds the relevant data.
- **Mandates testing of backups.** The policy identifies how often to test backups and the level of testing. For example, the policy may dictate performing full test restores weekly.
- **Identifies retention requirements.** How long data is held directly relates to how much backup media the organization must purchase and maintain. Laws or regulations may require retention of some data for several years and the organization can choose to limit retention of other data. Some organizations limit the amount of data they keep to reduce potential exposure to future legal proceedings. For example, a court order could direct administrators to comb through email for an investigation. The time spent will be significantly different if the organization kept archives of email only from the past year, or if it kept archives for the past 10 years.
- **Designates frequency of backups.** The business impact analysis (covered later in this chapter) helps an organization identify backup frequency by identifying recovery time objectives and recovery point objectives. This also helps determine the backup strategy, such as a full/incremental or full/differential strategy.
- **Protects backups.** Backup media is handled with the same level of protection as the original data. If an attacker gets a copy of a backup, it's a simple matter to restore it and access all the data. This helps prevent data loss or theft.
- **Identifies acceptable media disposal methods.** Backup media such as tapes holds a significant amount of information. Organizations often require the sanitation or destruction of tapes at the end of their life cycle. For example, you can erase all the data on a tape by degaussing it. A *degausser* is essentially a large magnet that makes the data unreadable. It's also possible to burn or shred tapes.

A key point here is that the backup policy identifies policy decisions related to backups. Administrators use the backup policy as a guide when creating backup plans.

Without a policy, these important decisions may never be addressed. The organization may not maintain off-site backups. They may not label backups. They may never test backups. All of this adds up to a catastrophe waiting to happen.

Remember this

Best practices associated with backups include storing a copy off-site for retention purposes, labeling the media, performing test restores, and destroying the media when it is no longer usable.

Comparing Business Continuity Elements

Business continuity planning helps an organization predict and plan for potential outages of critical services or functions. The goal is to ensure that critical business operations continue and the organization can survive the outage. Organizations often create a business continuity plan (BCP). This plan includes disaster recovery elements that provide the steps used to return critical functions to operation after an outage.

Disasters and outages can come from many sources, including:

- Fires
- Attacks
- Power outages
- Data loss from any cause
- Hardware and software failures
- Natural disasters, such as hurricanes, floods, tornadoes, and earthquakes

Addressing all of these possible sources takes a lot of time and effort. The goal is to predict the relevant disasters, their impact, and then develop recovery strategies to mitigate them. The overall process of business continuity planning generally takes the following steps:

1. Complete a business impact analysis.
2. Develop recovery strategies.
3. Develop recovery plans.
4. Test recovery plans.
5. Update plans.

You aren't required to know the details of all these steps for the CompTIA Security+ exam, but you should have a general idea of the process. The following sections provide an overview of the relevant topics.

Business Impact Analysis

A business impact analysis (BIA) is an important part of a BCP. It helps an organization identify critical systems and components that are essential to the organization's success. If critical systems and components fail and cannot be restored quickly, it's very possible that the organization will not survive the disaster.

For example, if a disaster such as a hurricane hit, what services must the organization restore to stay in business? Imagine a financial institution. It may decide that customers must have uninterrupted access to account data through an online site. If customers can't access their funds online, they may lose faith with the company and leave in droves.

On the other hand, the company may decide that it doesn't need the ability to accept and process loan applications right away. Loan processing is still important to the company's bottom line, but a delay will not seriously affect its ability to stay in business. In this case, the online site is a critical function, but the systems used for loan applications are not critical.

The time to make these decisions is not during a crisis. Instead, the organization completes a BIA in advance. The BIA involves collecting information from throughout the organization and documenting the results. This documentation identifies core business or mission requirements. The BIA does not recommend solutions. However, it provides management with valuable information so that they can focus on critical business functions. It helps them address some of the following questions:

- What are the critical systems and functions?
- Are there any dependencies related to these critical systems and functions?
- What is the maximum downtime limit of these critical systems and functions?
- What scenarios are most likely to impact these critical systems and functions?
- What is the potential loss from these scenarios?

As an example, imagine an organization earns an average of \$5,000 an hour through online sales. In this case, management might consider online sales to be a critical function and all systems that support online sales are critical systems. This includes web servers and back-end database servers. These servers depend on the network infrastructure connecting them, Internet access, and access to payment gateways for credit card charges.

After analysis, they might determine that the maximum allowable outage for online sales is five hours. Identifying the maximum downtime limit for the critical systems and functions is extremely important. It drives decisions related to recovery objectives and methods and helps an organization identify various contingency plans and policies.

The BIA evaluates various scenarios, such as fires, attacks, power outages, data loss, hardware and software failures, and natural disasters. Additionally, the BIA attempts to identify the potential loss from these scenarios. For example, a database server might host customer data, including credit card information. If an attacker was able to access this customer data, the cost to the organization might exceed millions of dollars.

You might remember the attack on retail giant Target during November and December 2013. Attackers accessed customer data on more than 110 million customers, resulting in significant losses for Target. Estimates of the total cost of the incident have ranged from \$600 million to over \$1 billion. This includes loss of sales—Target suffered a 46 percent drop in profits during the last quarter of 2013, compared with the previous year. Customers were afraid to use their credit cards in Target and simply stayed away. It also includes the cost to repair their image, the cost of purchasing credit monitoring for affected customers, fines from the payment-card industry, and an untold number of lawsuits. Target reportedly has \$100 million in cyber insurance that will help pay claims related to the data breach.

Remember this

The BIA identifies systems and components that are essential to the organization's success. It also identifies maximum downtime limits for these systems and components, various scenarios that can impact these systems and components, and the potential losses from an incident.

Recovery Time Objective

The *recovery time objective (RTO)* identifies the maximum amount of time it can take to restore a system after an outage. Many BIAs identify the maximum acceptable outage or maximum tolerable outage time for critical services or business functions. If an outage lasts longer than this maximum time, the impact is unacceptable to the organization.

For example, imagine an organization that sells products via a web site generates \$10,000 in revenue an hour. It might decide that the maximum acceptable outage for the web server is five minutes. With this in mind, the RTO is five minutes, indicating any outage must be limited to less than five minutes.

On the other hand, the organization may have a database server used by internal employees. Although the database server may be valuable, it is not critical. Management might decide they can accept an outage for as long as 24 hours, dictating an RTO less than 24 hours.

Recovery Point Objective

A recovery point objective (RPO) identifies a point in time where data loss is acceptable. As an example, a server may host archived data that has very few changes on a weekly basis. Management might decide that some data loss is acceptable, but they always want to be able to recover data from at least the previous week. In this case, the RPO is one week.

With an RPO of one week, administrators would ensure that they have at least weekly backups. In the event of a failure, they will be able to restore recent backups and meet the RPO.

In some cases, the RPO is up to the minute of the failure. For example, any data loss from an online database recording customer transactions might be unacceptable. In this case, the organization can use a variety of techniques to ensure administrators can restore data up to the moment of failure.

Remember this

The recovery time objective (RTO) identifies the maximum amount of time it should take to restore a system after an outage. It is derived from the maximum allowable outage time identified in the BIA. The recovery point objective (RPO) refers to the amount of data you can afford to lose.

Continuity of Operations

Continuity of operations planning (COOP) is an important element of a BCP. It focuses on restoring critical business functions at an alternate location after a critical outage. For example, if a hurricane or other disaster prevents the company from operating in one location, a COOP site allows it to continue to provide critical services at an alternate location. Many organizations plan for using a COOP site for as long as 30 days after relocating.

In this context, a site is an alternate location. It could be office space within a building, an entire building, or even a group of buildings. The four primary types of alternate sites are hot sites, cold sites, warm sites, and mobile sites.

As a rule of thumb, you'd use a hot site when you need to be operational within 60 minutes, or a cold site if you must be operational within a few days. For periods between these two extremes, you'd use a warm site. A mobile site is an alternate location when an organization doesn't want to have a permanent location as an alternate site. The following sections provide more details on these sites.

Remember this

Continuity of operations planning (COOP) sites provide an alternate location for operations after a critical outage. The most common sites are hot, cold, warm, and mobile sites.

Hot Site

A hot site would be up and operational 24 hours a day, seven days a week and would be able to take over functionality from the primary site quickly after a primary site failure. It would include all the equipment, software, and communication capabilities of the primary site, and all the data would be up to date. In many cases, copies of backup tapes are stored at the hot site as the off-site location.

In many cases, a hot site is another active business location that has the capability to assume operations during a disaster. For example, a financial institution could have locations in two separate cities. The second location provides noncritical support services, but also includes all the resources necessary to assume the functions of the first location.

Some definitions of hot sites indicate they can take over instantaneously, though this isn't consistent. In most cases, it takes a little bit of time to transfer operations to the hot site, and this can take anywhere from a few minutes to an hour.

Clearly, a hot site is the most effective disaster recovery solution for high-availability requirements. If an organization must keep critical systems with high-availability requirements, the hot site is the best choice. However, a hot site is the most expensive to maintain and keep up to date.

Remember this

A hot site includes personnel, equipment, software, and communication capabilities of the primary site with all the data up to date. A hot site provides the shortest recovery time compared with warm and cold sites. It is the most effective disaster recovery solution but is also the most expensive to maintain.

Cold Site

A cold site requires power and connectivity but not much else. Generally, if it has a roof, electricity, running water, and Internet access, you're good to go. The organization brings all the equipment, software, and data to the site when it activates it.

I often take my dogs for a walk at a local army base and occasionally see soldiers activate an extreme example of a cold site. On most weekends, the fields are empty. Other weekends, soldiers have transformed one or more fields into complete operational sites with tents, antennas, cables, generators, and porta-potties.

Because the army has several buildings on the base, they don't need to operate in the middle of fields, but what they're really doing is testing their ability to stand up a cold site wherever they want. If they can do it in the field, they can do it in the middle of a desert, or anywhere else they need to.

A cold site is the cheapest to maintain, but it is also the most difficult to test.

Warm Site

You can think of a warm site as the Goldilocks solution—not too hot and not too cold, but just right. Hot sites are generally too expensive for most organizations, and cold sites generally take too long to configure for full operation. However, the warm site provides a compromise that an organization can tailor to meet its needs.

For example, an organization can place all the necessary hardware at the warm site location but not include up-to-date data. If a disaster occurs, the organization can copy the data to the warm site and take over operations. This is only one example, but there are many different possibilities of warm site configurations.

Site Variations

Although hot, cold, and warm sites are the most common, you may also come across two additional alternate site types: mobile and mirrored.

A *mobile site* is a self-contained transportable unit with all the equipment needed for specific requirements. For example, you can outfit a semitrailer with everything needed for operations, including a satellite dish for connectivity. Trucks, trains, or ships haul it to its destination and it only needs power to start operating.

Mirrored sites are identical to the primary location and provide 100 percent availability. They use real-time transfers to send modifications from the primary location to the mirrored site. Although a hot site can be up and operational within an hour, the mirrored site is always up and operational.

Remember this

A cold site will have power and connectivity needed for COOP activation, but little else. Cold sites are the least expensive and the hardest to test. A warm site is a compromise between a hot site and a cold site. Mobile sites do not have dedicated locations but can provide temporary support during a disaster.

After the Disaster

After the disaster has passed, you will want to return all the functions to the primary site. As a best practice, organizations return the least critical functions to the primary site first. Remember, the critical functions are operational at the alternate site and can stay there as long as necessary.

If a site has just gone through a disaster, it’s very likely that there are still some unknown problems. By enabling the least critical functions first, undiscovered problems will appear and can be resolved without significantly affecting critical business functions.

Disaster Recovery

Disaster recovery is a part of an overall business continuity plan. Often the organization will use the business impact analysis to identify the critical systems and components and then develop disaster recovery strategies and disaster recovery plans (DRPs) to address the systems hosting these functions.

In some cases, an organization will have multiple DRPs within a BCP, and in other cases, the organization will have a single DRP. For example, it's possible to have individual DRPs that identify the steps to recover individual critical servers, and other DRPs that detail the recovery steps after different types of disasters such as hurricanes or tornadoes. A smaller organization may have a single DRP that simply identifies all the steps used to respond to any disruption.

A DRP or a BCP will include a hierarchical list of critical systems. This list identifies what systems to restore after a disaster and in what order. For example, should a server hosting an online web site be restored first, or a server hosting an internal application? The answer is dependent on how the organization values and uses these servers. In some cases, systems have interdependencies requiring systems to be restored in a certain order.

If the DRP doesn't prioritize the systems, individuals restoring the systems will use their own judgment, which may not meet the overall needs of the organization. For example, Nicky New Guy may not realize that a web server is generating \$5,000 an hour in revenue but does know that he's responsible for keeping a generic file server operational. Without an ordered list of critical systems, he may spend his time restoring the file server and not the web server.

This hierarchical list is valuable when using alternate sites such as warm or cold sites, too. When the organization needs to move operations to an alternate site, the organization will want the most important systems and functions restored first.

Similarly, the DRP often prioritizes the services to restore after an outage. As a rule, critical business functions and security services are restored first. Support services are restored last.

The different phases of a disaster recovery process typically include the following steps:

- **Activate the disaster recovery plan.** Some disasters, such as earthquakes or tornadoes, occur without much warning, and a disaster recovery plan is activated after the disaster. Other disasters, such as hurricanes, provide a warning, and the plan is activated when the disaster is imminent.
- **Implement contingencies.** If the recovery plan requires implementation of an alternate site, critical functions are moved to these sites. If the disaster destroyed on-site backups, this step retrieves the off-site backups from the off-site location.
- **Recover critical systems.** After the disaster has passed, the organization begins recovering critical systems. The DRP documents which systems to recover and includes detailed steps on how to recover them. This also includes reviewing change management documentation to ensure that recovered systems include approved changes.
- **Test recovered systems.** Before bringing systems online, administrators test and verify them. This may include comparing the restored system with a performance baseline to verify functionality.
- **Document and review.** The final phase of disaster recovery includes a review of the disaster, sometimes called an after-action review. This often includes a lessons-learned review to identify what went right and what went wrong. The organization often updates the plan after a disaster to incorporate any lessons learned.

Remember this

A disaster recovery plan (DRP) includes a hierarchical list of critical systems and often prioritizes services to restore after an outage. Testing validates the plan. The final phase of disaster recovery includes a review to identify any lessons learned and may include an update of the plan.

Planning for Communications

When planning for any disaster or major disruption, it's important to plan for communications. Normal communications methods might not be available during an incident. For example, personnel might commonly communicate via email, but email services might be down during an emergency. Similarly, cell-based phone lines might not be operational.

One alternate method of communications is the use of a war room. This can be as simple as a conference room, but transformed into a central command center. Response team members report on their progress to people in the war room, and if anyone needs to get up-to-date information, they go to the war room to get it.

A communication plan will include methods of communicating with the following entities:

- **Disaster response team members.** This might be to inform them of an impending disaster or to keep in touch during the recovery. Some organizations use push-to-talk phones that aren't affected by cell-phone outages.
- **Employees.** It's common for a plan to require mission-essential employees to come to work, but tell non-mission-essential personnel to stay home. Many organizations coordinate with TV and radio stations to advertise their decisions during many weather disasters such as hurricanes.
- **Customers.** Many organizations let customers know that they are responding to a disaster via a web page. This is especially useful if the disaster prevents the organization from providing certain services. As an example, an online banking site might post a notification stating that the bank is implementing a disaster recovery plan in response to an emergency. It would include information on when they expect services to return. Customers will understand this much better than going to their online bank and simply finding that it is down.
- **Suppliers.** In some cases, suppliers might need to halt deliveries for a period of time. By calling them before the delivery, it prevents unnecessary problems.
- **Media.** Whenever possible, it's best to defer all media requests to a public relations (PR) expert within the organization. They know how to project the right image to the media to prevent miscommunication problems. For example, a TV reporter might ask a technician for a comment and he might reply, "We got slammed! It's chaos in there!" In contrast, a PR expert might say something like, "This disaster hit us hard. However, we are putting our recovery plans into action and facing this disaster with all available resources." The second response projects much more confidence than the first. The communication plan can also include templates to respond to media requests. For example, a template might include, "We are putting our recovery plans into action and facing this disaster with all available resources. We expect to have more information on the impact soon."
- **Regulatory agencies.** Some organizations must report certain events to regulatory agencies. For example, if an attack results in a data breach of customer Personally Identifiable Information (PII), the organization might have a legal obligation to report it. It's important to document what must be reported, and how to do so.

Remember this

BCPs and DRPs commonly include a communication plan. It identifies alternate methods of communication, such as a war room or push-to-talk phones. It also identifies who must be contacted, such as response team members, employees, suppliers, customers, media, and regulatory agencies.

IT Contingency Planning

Information technology (IT) contingency planning is focused on recovery for IT systems only. From a broader perspective, a BCP looks at the entire organization and can include one or more DRPs. A DRP provides steps and procedures to return one or more systems to operation after a major disruption or outage and may involve moving operations to a different location. IT contingency planning works on a smaller scale and examines single systems only.

Notice there is a little overlap here. A DRP can document steps and procedures to return a single system to operation, just as an IT contingency plan can. For a small organization, there isn't any real distinction between the two. However, for a large organization, the IT contingency plan provides a little more manageability for the process. Instead of creating a massive DRP that does everything, it can create a DRP with multiple IT contingency plans.

Succession Planning

Succession planning means different things depending on the context. From a purely business perspective, succession planning identifies people within the organization who can fill leadership positions. The goal is to ensure the business can continue to thrive even if key leaders within the organization become ill or leave unexpectedly.

Within the context of business continuity and disaster preparedness, succession planning is the process of identifying a hierarchical chain of command. For example, the cost to activate a warm site might be substantial, so a BCP may dictate that only the CEO can decide when to activate the warm site. However, what does the organization do if the CEO is vacationing in Las Vegas or is otherwise unreachable? The BCP can provide an order of succession used during a disaster to ensure that someone has the authority to make decisions.

Without clear succession planning, people might try to take control when they shouldn't. Even when succession planning is clear, people might still try to take control in crisis situations. In 1981, after John Hinckley Jr. shot President Reagan, Secretary of State Alexander Haig infamously said, "I am in control here." However, the U.S. constitution identifies the line of succession if the president is incapacitated. The succession order is the vice president, the Speaker of the House, the president pro tempore of the Senate, and then the secretary of state. Several people quickly pointed out the correct line of succession.

In contrast, if no one perceives they have any authority, they might be reluctant to make any decisions. To address this, a BCP or DRP might include a chart of roles and responsibilities organized by title, along with a hierarchical chain of command. This makes clear to personnel who can make certain decisions.

Remember this

Succession planning ensures that an organization can continue to thrive even if key leaders unexpectedly leave or are unavailable. Some BCPs and DRPs include charts identifying roles and responsibilities, along with a clear chain of command.

BCP and DRP Testing

Business continuity plans and disaster recovery plans include testing. Testing validates that the plan works as desired and will often include testing redundancies and backups. There are several different types of testing used with BCPs and DRPs.

National Institute of Standards and Technology (NIST) Special Publication (SP) 800-34, “Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities,” provides detailed guidance on testing BCP and DRP plans. SP 800-34 identifies two primary types of exercises: tabletop exercises and functional exercises.

A *tabletop exercise* (also called a desktop exercise or a structured walk-through) is discussion-based. A coordinator gathers participants in a classroom or conference room, and leads them through one or more scenarios. As the coordinator introduces each stage of an incident, the participants identify what they’ll do based on the plan. This generates discussion about team members’ roles and responsibilities, and the decision-making process during an incident. Ideally, this validates that the plan is valid. However, it sometimes reveals flaws. The BCP coordinator ensures the plans are rewritten if necessary.

Functional exercises provide personnel with an opportunity to test the plans in a simulated operational environment. There is a wide range of functional exercises, from simple simulations to full-blown tests. In a simulation, the participants go through the steps in a controlled manner without affecting the actual system. For example, a simulation can start by indicating that a server failed. Participants then follow the steps to rebuild the server on a test system. A full-blown test goes through all the steps of the plan. In addition to verifying that the test works, this also shows the amount of time it will take to execute the plan.

Some of the common elements of testing include:

- **Backups.** Backups are tested by restoring the data from the backup, as discussed in the “Testing Backups” section earlier in this chapter.
- **Server restoration.** A simple disaster recovery exercise rebuilds a server. Participants follow the steps to rebuild a server using a test system without touching the live system.
- **Server redundancy.** If a server is within a failover cluster, you can test the cluster by taking a primary node offline. Another node within the cluster should automatically assume the role of this offline node.
- **Alternate sites.** You can test an alternate site (hot, cold, or warm) by moving some of the functionality to the alternate site and ensuring the alternate site works as desired. It’s also possible to test individual elements of an alternate site, such as Internet connectivity, or the ability to obtain and restore backup media.

Remember this

You can validate business continuity plans and disaster recovery plans through testing. Tabletop exercises are discussion-based only and are typically performed in a classroom or conference setting. Functional exercises are hands-on exercises.

Testing Controls

In some plans, it’s appropriate to test specific security controls to see if they operate as desired. For example, if the primary node in a failover cluster fails, it should automatically failover to the secondary node. An extreme method of testing this is to turn off the primary node and see if the secondary node takes over.

In some cases, an organization might want to ensure that employees can continue mission-essential operations even if certain resources are unavailable. For example, a DRP might identify alternate methods the HR department can use to perform certain functions even if the HR server is unavailable. By turning off the HR server for a day, the organization can determine if the HR personnel are able to perform their jobs using these alternate methods.

Escape Plans, Escape Routes, and Drills

Safety of personnel is always a concern and some BCP and DRPs include escape plans, letting people know how to escape and the route to take. As a simple example, rooms might include a map showing the shortest route out of a building in case of fire. Similarly, some cities have escape routes people can take in the event of a major storm such as a hurricane.

When escape plans and escape routes are in place, it’s useful to perform drills periodically to test the plans and routes. Many organizations do routine fire drills to ensure that personnel are able to evacuate a building in a short period of time.

Implementing Environmental Controls

Although environmental controls might not seem security related, they directly contribute to the availability of systems. This includes ensuring temperature and humidity controls are operating properly, fire suppression systems are in place, and proper procedures are used when running cables.

Heating, Ventilation, and Air Conditioning

Heating, ventilation, and air conditioning (HVAC) systems are important physical security controls that enhance the availability of systems. Quite simply, computers and other electronic equipment can't handle drastic changes in temperatures, especially hot temperatures. If systems overheat, the chips can actually burn themselves out.

The cooling capacity of HVAC systems is measured as tonnage. This has nothing to do with weight, but instead refers to cooling capacity. One ton of cooling equals 12,000 British thermal units per hour (Btu/hour), and typical home HVAC systems are three-ton units. Higher-tonnage HVAC systems can cool larger areas or areas with equipment generating more heat.

The amount of air conditioning needed to cool a massive data center is much greater than you need to cool your home, primarily because of all the heat generated by the equipment. If your home air conditioner fails in the middle of summer, you may be a little uncomfortable for a while, but if the data center HVAC system fails, it can result in loss of availability and a substantial loss of money.

Chapter 8, "Managing Risk," mentions mean time between failures (MTBF) as a measure of a system's reliability. If temperatures aren't controlled within a data center, it decreases the MTBF times, resulting in more frequent failures. In contrast, if an HVAC system keeps systems at a consistent temperature, it tends to increase the MTBF times, resulting in fewer failures.

I worked in several environments where we had a policy of shutting down all electronics when the room temperature reached a certain threshold. When we didn't follow the policy, the systems often developed problems due to the heat and ended up out of commission for a lot longer than the AC.

Most servers aren't in cases like a typical desktop computer. Instead, they are housed in rack-mountable cases. These rack-mountable servers are installed in equipment cabinets (also called racks or bays) about the size of tall refrigerators. A large data center will have multiple cabinets lined up beside each other in multiple rows.

These cabinets usually have locking doors in the front and rear for physical security. The doors are perforated with cold air coming in the front, passing over and through the servers to keep them cool, and warmer air exiting out the rear. Additionally, a server room has raised flooring with air conditioning pumping through the space under the raised floor.

Remember this

Higher-tonnage HVAC systems provide more cooling capacity. This keeps server rooms at lower operating temperatures, and results in fewer failures and longer MTBF times.

Hot and Cold Aisles

Hot and cold aisles help regulate the cooling in data centers with multiple rows of cabinets. The back of all the cabinets in one row will face the back of all the cabinets in an adjacent row. Because the hot air exits out the back of the cabinet, the aisle with the backs facing each other is the hot aisle.

Similarly, the front of the cabinets in one row is facing the front of the cabinets in the adjacent row. Cool air is pumped through the floor to this cool aisle using perforated floor tiles in the raised flooring. This is the cold aisle. In some designs, cool air is also pumped through the base of the cabinets. This depends on the design of the cabinets and the needs of the equipment.

Consider what happens if all the cabinets had their front facing the same way without a hot/cold aisle design. The hot air pumping out the back of one row of cabinets would be sent to the front of the cabinets behind them. The front row would have very cold air coming in the front, but other rows would have warmer air coming in the front.

Of course, an HVAC also includes a thermostat as a temperature control and additional humidity controls. The thermostat ensures that the air temperature is controlled and maintained. Similarly, humidity controls ensure that the humidity is controlled. High humidity can cause condensation on the equipment, which causes water damage. Low humidity allows a higher incidence of electrostatic discharge (ESD).

HVAC and Fire

HVAC systems are often integrated with fire alarm systems to help prevent a fire from spreading. One of the core elements of a fire is oxygen. If the HVAC system continues to operate normally while a fire is active, it continues to pump oxygen, which feeds the fire. When the HVAC system is integrated with the fire alarm system, it controls the airflow to help prevent the rapid spread of the fire. Many current HVAC systems have dampers that can control airflow to specific areas of a building. Other HVAC systems automatically turn off when fire suppression systems detect a fire.

Remember this

HVAC systems increase availability by controlling temperature and humidity. Temperature controls help ensure a relatively constant temperature. Humidity controls reduce the potential for damage from electrostatic discharge and damage from condensation. HVAC systems should be integrated with the fire alarm systems and either have dampers or the ability to be turned off in the event of a fire.

Fail-Safe Versus Fail-Open

Many times, it's important to consider the state of a system if it fails. You can often force a system to fail in an open state, or to fail in a safe or secure state. The terms fail-safe, fail-secure, and fail-close all mean the same thing. The state you choose is often dependent on the needs of the organization.

For example, you may have a system that requires high availability, but security isn't as important. If it fails, you would want it to fail in an open state so that it remains available.

Consider an exit door secured with a proximity card. Normally, employees open the door with the proximity card and the system records their exit. The proximity card provides security, but you need the exit to remain highly available.

What happens if a fire starts and power to the building is lost? The proximity card reader won't work, and if the door can't open, employees will be trapped. In this case, you would want the proximity card reader to fail in the fail-open state so that personnel can get out. The value of personnel safety is always paramount. Of course, this does introduce a vulnerability. An attacker might be able to access a secure data center by destroying a proximity card reader.

On the other hand, consider a firewall used to provide security for a network. In this case, security is more important than availability. For example, if a firewall access control list (ACL) became corrupt, you would want it to fail in a fail-safe or secure mode. Essentially, it would block all traffic and continue to provide security for the network.

Fire Suppression

You can fight fires with individual fire extinguishers, with fixed systems, or both. Most organizations included fixed systems to control fires and place portable fire extinguishers in different areas around the organization. A fixed system can detect a fire and automatically activate to extinguish the fire. Individuals use portable fire extinguishers to suppress small fires.

The different components of a fire are heat, oxygen, fuel, and a chain reaction creating the fire. Fire suppression methods attempt to remove or disrupt one of these elements to extinguish a fire. You can extinguish a fire using one of these methods:

- **Remove the heat.** Fire extinguishers commonly use chemical agents or water to remove the heat. However, water should never be used on an electrical fire.
- **Remove the oxygen.** Many methods use a gas, such as carbon dioxide (CO₂) to displace the oxygen. This is a common method of fighting electrical fires because CO₂ and similar gasses are harmless to electrical equipment.
- **Remove the fuel.** Fire-suppression methods don't typically fight a fire this way, but of course, the fire will go out once all the material is burned.
- **Disrupt the chain reaction.** Some chemicals can disrupt the chain reaction of fires to stop them.

The class of fire often determines what element of the fire you will try to remove or disrupt. Within the United States, fires are categorized in one of the following fire classes:

- **Class A—Ordinary combustibles.** These include wood, paper, cloth, rubber, trash, and plastics.
- **Class B—Flammable liquids.** These include gasoline, propane, solvents, oil, paint, lacquers, and other synthetics or oil-based products.
- **Class C—Electrical equipment.** This includes computers, wiring, controls, motors, and appliances. The CompTIA Security+ exam is computer-centric, so you should especially understand that a Class C fire is from electrical equipment. You should not fight Class C fires with water or water-based materials, such as foam, because the water is conductive and can pose significant risks to personnel.
- **Class D—Combustible metals.** This includes metals such as magnesium, lithium, titanium, and sodium. Once they start to burn, they are much more difficult to extinguish than other materials.

You can extinguish a Class A fire with water to remove the heat. However, water makes things much worse if you use it on any of the other classes. For example, using water on live equipment actually poses a risk because electricity can travel up the water stream and shock you. Additionally, water damages electrical equipment.

Environmental Monitoring

Environmental monitoring includes temperature and humidity controls. From a very basic perspective, an HVAC system monitors the current temperature and humidity and makes adjustments as necessary to keep the temperature and humidity constant.

Large-scale data centers often have sophisticated logging capabilities for environmental monitoring. The HVAC system still attempts to keep the temperature and humidity constant. However, the logs record the actual temperature and humidity at different times during the day. This allows administrators to review the performance of the HVAC system, to see if it is able to keep up with the demands within the data center.

Shielding

Shielding helps prevent electromagnetic interference (EMI) and radio frequency interference (RFI) from interfering with normal signal transmissions. It also protects against unwanted emissions and helps prevent an attacker from capturing network traffic.

Although you may see EMI and RFI in the same category as EMI/RFI, they are different. EMI comes from different types of motors, power lines, and even fluorescent lights. RFI comes from radio frequency (RF) sources such as AM or FM transmitters. However, shielding used to block interference from both EMI and RFI sources is often referred to as simply EMI shielding.

Attackers often use different types of eavesdropping methods to capture network traffic. If the data is emanating outside of the wire or outside of an enclosure, attackers may be able to capture and read the data. EMI shielding fulfills the dual purpose of keeping interference out and preventing attackers from capturing network traffic.

Shielding Cables

Twisted-pair cable, such as CAT5e and CAT6 cable, comes in both shielded twisted-pair (STP) and unshielded twisted-pair (UTP) versions. The shielding helps prevent an attacker from capturing network traffic and helps block interference from corrupting the data.

When data travels along a copper wire (such as twisted-pair), it creates an induction field around the wire. If you have the right tools, you can simply place the tool around the wire and capture the signal. The shielding in STP cable blocks this. Fiber-optic cable is not susceptible to this type of attack. Signals travel along a fiber-optic cable as light pulses, and they do not create an induction field.

Remember this

EMI shielding prevents outside interference sources from corrupting data and prevents data from emanating outside the cable.

Protected Distribution of Cabling

Physical security includes planning where you route cables and how you route them. Skilled network administrators can cut a twisted-pair cable, attach an RJ-45 connector to each end, and connect them back together with an adapter in less than 5 minutes. I recently taught a CompTIA Security+ class to some Verizon Fiber Optic Services FiOS technicians, and they said they can do the same thing with a fiber-optic cable within 10 minutes.

If an attacker did this, he could connect the cut cable with a hub, and then capture all the traffic going through the hub with a protocol analyzer. This represents a significant risk.

One method of reducing this risk is to run cables through cable troughs. A *cable trough* is a long metal container, typically about 4 inches wide by 4 inches high. If you run data cables through the cable trough, they aren't as accessible to potential attackers. In contrast, many organizations simply run the cable through a false ceiling or a raised floor.

In addition to considering physical security, it's important to keep the cables away from EMI sources. As an example, if technicians run cables over or through fluorescent lighting fixtures, the EMI from the lights can disrupt the signals on the cables. The result is intermittent connectivity for users.

Faraday Cage

A *Faraday cage* is a room that prevents signals from emanating beyond the room. It includes electrical features that cause RF signals that reach the boundary of the room to be reflected back, preventing signal emanation outside the Faraday cage.

In addition to preventing signals from emanating outside the room, a Faraday cage also provides shielding to prevent outside interference such as EMI and RFI from entering the room.

At a very basic level, some elevators act as a Faraday cage (though I seriously doubt the designers were striving to do so). You may have stepped into an elevator and found that your cell phone stopped receiving and transmitting signals. The metal shielding around the elevator prevents signals from emanating out or signals such as the cell phone tower signal from entering the elevator.

On a smaller scale, electrical devices such as computers include shielding to prevent signals from emanating out and block interference from getting in.

Chapter 9 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Adding Redundancy

- A single point of failure is any component that can cause the entire system to fail if it fails.
- RAID disk subsystems provide fault tolerance and increase availability. RAID-1 (mirroring) uses two disks, RAID-5 uses three or more disks and can survive the failure of one disk, and RAID-6 uses four or more disks and can survive the failure of two disks.
- Server redundancies include failover clusters and load balancing. Failover clusters remove a server as a single point of failure. If one node in a cluster fails, another node can take over.
- Load balancing spreads the processing load over multiple servers to ensure availability when the processing load increases. Many web-based applications use load balancing for higher availability.
- A UPS system provides fault tolerance for power fluctuations and provides short-term power for systems during power outages. Generators provide long-term power for systems during extended power outages.

Protecting Data with Backups

- Backup strategies include full, full/differential, and full/incremental strategies. A full backup strategy alone allows the quickest recovery time.
- Full/incremental backup strategies minimize the amount of time needed to perform daily backups.
- Test restores verify the integrity of backups. A test restore of a full backup verifies a backup can be restored in its entirety.
- Backups should be labeled to identify the contents. A copy of backups should be kept off-site.

Comparing Business Continuity Elements

- A business impact analysis (BIA) is part of a business continuity plan (BCP) and it identifies systems and components that are essential to the organization's success.
- The BIA identifies maximum downtimes for these systems and components, various scenarios that can affect these systems and components, and the potential losses from an incident.
- Recovery time objective (RTO) identifies the maximum amount of time it should take to restore a system after an outage. The recovery point objective (RPO) refers to the amount of data you can afford to lose.
- Continuity of operations planning (COOP) sites provide alternate locations for business functions after a major disaster.
- A hot site includes everything needed to be operational within 60 minutes. It is the most effective recovery solution and the most expensive. A cold site has power and connectivity requirements and little else. It is the least expensive to maintain.
- Warm sites are a compromise between hot sites and cold sites. Mobile sites do not have dedicated locations, but can provide temporary support during a disaster.
- Disaster recovery planning is part of overall business continuity planning. A disaster recovery plan (DRP) includes the steps to return one or more systems to full operation. BCPs or DRPs include a hierarchical list of critical systems identifying the order of restoration.
- BCPs and DRPs commonly include a communication plan. It identifies alternate methods of communication, such as a war room or push-to-talk phones. It also identifies who to contact, such as response team members, employees, suppliers, customers, media, and regulatory agencies.
- Succession planning ensures that an organization can continue to operate even if key leaders are unavailable. Succession planning charts identify roles and responsibilities to follow during a disaster, along with a clear chain of command.
- Periodic testing validates BCPs and DRPs. Disaster recovery exercises validate the steps to restore individual systems, activate alternate sites, and other actions documented within a plan. Tabletop exercises are discussion-based only. Functional exercises are hands-on exercises.

Implementing Environmental Controls

- Heating, ventilation, and air conditioning (HVAC) systems control airflow for data centers and server rooms. Temperature controls protect systems from damage due to overheating.
- Higher-tonnage HVAC systems provide more cooling capacity. You can increase the mean time between failures (MTBF) times and overall availability by keeping server rooms at lower operating temperatures.
- Humidity controls protect against ESD damage by ensuring humidity isn't too low. They also protect against water damage from condensation if humidity gets too high.
- HVAC systems should be integrated with the fire alarm systems and either have dampers or the ability to be turned off in the event of a fire.
- EMI shielding prevents problems from EMI sources such as fluorescent lighting fixtures. It also prevents data loss in twisted-pair cables.

Chapter 9 Practice Questions

1. An organization needs to improve fault tolerance to increase data availability. However, the organization has a limited budget. Which of the following is the BEST choice to meet the organization's needs?
 - A. RAID
 - B. Backup system
 - C. Cluster
 - D. UPS
2. Your organization hosts a web site with a back-end database server. During a recent power outage, the server crashed, resulting in a significant amount of lost data. Which of the following can your organization implement to prevent this loss from occurring again?
 - A. Redundancy
 - B. Disaster recovery procedures
 - C. Warm site
 - D. Higher RTO
3. A network administrator configured several servers to work together to increase the processing capabilities for a web application. What does the administrator MOST likely implement?
 - A. Failover clustering
 - B. RAID-6
 - C. EMI shielding
 - D. Load balancing
4. Your company's web site experiences a large number of client requests during certain times of the year. Which of the following could your company add to ensure the web site's availability during these times?
 - A. Fail-open cluster
 - B. Certificates
 - C. Web application firewall
 - D. Load balancing
5. Your organization hosts a high-volume web site, which generates a significant amount of revenue. You are asked to recommend a method to increase the availability of this web site. Which of the following choices is the BEST choice?
 - A. Load balancing
 - B. Hot site
 - C. WAF
 - D. UTM
6. Your backup policy for a database server dictates that the amount of time needed to perform backups should be minimized. Which of the following backup plans would BEST meet this need?
 - A. Full backups on Sunday and full backups every other day of the week
 - B. Full backups on Sunday and differential backups every other day of the week
 - C. Full backups on Sunday and incremental backups every other day of the week
 - D. Differential backups on Sunday and incremental backups every other day of the week
7. A business continuity expert is creating a BIA. Which of the following elements is MOST likely to be omitted from the BIA?
 - A. List of critical systems and functions
 - B. Recommended solutions
 - C. Critical downtime limit
 - D. Potential loss
8. After a recent attack causing a data breach, an executive is analyzing the financial losses. She determined that the attack is likely to cost at least \$1

million. She wants to ensure that this information is documented for future planning purposes. Where is she MOST likely to document it?

- A. DRP
- B. BIA
- C. COOP
- D. RTO

9. You are helping implement your company’s business continuity plan. For one system, the plan requires an RTO of five hours and an RPO of one day. Which of the following would meet this requirement?

- A. Ensure the system can be restored within five hours and ensure it does not lose more than one day of data.
- B. Ensure the system can be restored within one day and ensure it does not lose more than five hours of data.
- C. Ensure the system can be restored between five hours and one day after an outage.
- D. Ensure critical systems can be restored within five hours and noncritical systems can be restored within one day.

10. An organization is considering an alternate location as part of its business continuity plan. It wants to identify a solution that provides the shortest recovery time. What will it choose?

- A. Cold site
- B. Warm site
- C. Hot site
- D. Succession site

11. Your organization is working on its business continuity plan. Management wants to ensure that documents provide detailed information on what technicians should do after an outage. Specifically, they want to list the systems to restore and the order in which to restore them. What document includes this information?

- A. HVAC
- B. BIA
- C. DRP
- D. Succession plan

12. Your organization is updating its disaster recovery documents. You’re asked to review the communication plans for possible updates. Which of the following should you ensure is included in the communication plan?

- A. A list of test plans and procedures
- B. The succession plan
- C. Methods used to communicate with response team members, employees, suppliers, and customers
- D. List of scenarios with potential loss statements

13. A BCP includes a chart listing roles within the organization along with their matching responsibilities during a disaster. It also includes a chain of command. What is the purpose of this chart?

- A. IT contingency planning
- B. Succession planning
- C. COOP
- D. RTO

14. The BCP coordinator at your organization is leading a meeting on-site with key disaster recovery personnel. The purpose of the meeting is to perform a test. What type of test is this?

- A. Functional exercise
- B. Full-blown test
- C. Tabletop exercise
- D. Simulation to perform steps of a plan

15. Personnel within your organization turned off the HR data server for over six hours to perform a test. Which of the following is the MOST likely purpose of this?

- A. BIA

- B. Succession planning
- C. Tabletop exercises
- D. COOP

16. Humidity controls in your data center are failing. You need to convince management of the importance of these. What would you tell them?

- A. Failing humidity controls can cause damage from EMI and ESD.
- B. Failing humidity controls can cause damage from temperature variations and EMI.
- C. Failing humidity controls can cause damage from condensation and poor ventilation.
- D. Failing humidity controls can cause damage from ESD and condensation.

17. Your organization is evaluating replacement HVAC systems and is considering increasing current capacities. Which of the following is a potential security benefit of increasing the HVAC capabilities?

- A. Lower MTBF times of hardware components due to lower temperatures
- B. Higher MTBF times of hardware components due to lower temperatures
- C. Lower MTTR times of hardware components due to lower temperatures
- D. Higher MTTR times of hardware components due to lower temperatures

18. Without adequate physical security controls, attackers can cause significant damage to systems within a data center. Which of the following could an attacker manipulate to cause extensive physical damage?

- A. Video surveillance systems
- B. Environmental controls
- C. Firewall ACLs
- D. IDS settings

19. An attacker was able to sneak into your building but was unable to open the server room door. He bashed the proximity badge reader with a portable fire extinguisher and the door opened. What is the MOST likely reason that the door opened?

- A. The access system was designed to fail-open.
- B. The access system was designed to fail-close.
- C. The access system was improperly installed.
- D. The portable fire extinguisher included a proximity badge.

20. Which of the following is an environmental control?

- A. EMI shielding
- B. Fencing
- C. Video surveillance
- D. Motion detection

Chapter 9 Practice Question Answers

1. **A.** A redundant array of inexpensive disks (RAID) system would provide fault tolerance for disk drives and increase data availability if drives fail. A backup system improves data availability because you can restore data after data is lost or corrupt. However, a backup system does not provide fault tolerance. A cluster provides fault tolerance at the server level and ensures a service continues to operate even if a server fails. However, a cluster is more expensive than a RAID. An uninterruptible power supply (UPS) provides short-term power after a power failure but does not directly increase data availability.

2. **A.** Server redundancy solutions such as a failover cluster would prevent this type of loss. Additionally, a power redundancy solution such as an uninterruptible power supply (UPS) would prevent this. Disaster recovery procedures help restore the systems after a disaster, but they wouldn't prevent the incident. A warm site is as an alternate site, but it wouldn't prevent data loss. The recovery time objective (RTO) identifies the time period when you plan to restore a system after an outage, but it doesn't prevent a loss.

3. **D.** The administrator most likely implemented servers to work together in a load-balancing configuration. Load balancing shifts the load between multiple servers to increase the number of clients the application can handle, ultimately increasing the overall processing capabilities. Failover clustering adds one or more servers for high availability and a redundant array of inexpensive disks 6 (RAID-6) provides fault tolerance for the disk subsystem, but neither increases processing capabilities. Electromagnetic interference (EMI) shielding is an environmental control that can protect

against intermittent problems due to EMI.

4. **D.** Load balancing shifts the load among multiple systems and can increase the site's availability by adding additional nodes when necessary. A failover cluster also provides high availability, but there is no such thing as a fail-open cluster. Certificates help ensure confidentiality and integrity, but do not assist with availability. A web application firewall helps protect a web server against attacks, but it does not increase availability from normal client requests.
5. **A.** Load balancing adds additional servers to a service and shares the load among the servers. This increases availability because a single server is not overloaded and additional servers can be added as needed. A hot site supports operations at an alternate site after a disaster, but it is very expensive and not the best choice if you only want to increase the availability of a web site. A web application firewall (WAF) and a unified threat management (UTM) device both provide security, but they do not directly address availability.
6. **C.** A full/incremental backup strategy is best with one full backup on one day and incremental backups on the other days. A full backup every day would require the most time every day. Differential backups become steadily larger as the week progresses and take more time to back up than incremental backups. Backups must start with a full backup, so a differential/incremental backup strategy is not possible.
7. **B.** A business impact analysis (BIA) does not include recommended solutions. It does identify critical systems and functions, dependencies, critical downtime limits, potential scenarios causing a loss, and the potential loss.
8. **B.** A business impact analysis (BIA) includes information on potential losses and is the most likely document of those listed where this loss would be documented. A disaster recovery plan (DRP) includes methods used to recover from an outage. Continuity of operations planning (COOP) includes methods, such as alternate sites, used to keep an organization operational after an outage. The recovery time objective (RTO) identifies the time period when you plan to restore a system after an outage; it is not a document.
9. **A.** The recovery time objective (RTO) identifies the maximum amount of time it should take to restore a system after an outage. The recovery point objective (RPO) refers to the amount of data you can afford to lose. RTO only refers to time, not data. RPO refers to data recovery points, not time to restore a system.
10. **C.** A hot site has the shortest recovery time, but it is also the most expensive. Cold sites have the longest recovery time, and warm sites are shorter than cold sites but not as quick as hot sites. Succession site isn't a valid type of alternate location.
11. **C.** The disaster recovery plan (DRP) typically includes a hierarchical list of critical systems that identifies what to restore and in what order. Heating, ventilation, and air conditioning (HVAC) is not a document. The business impact analysis (BIA) identifies critical systems and components but does not include recovery methods or procedures. Succession planning refers to people, not systems, and it clarifies who can make decisions during a disaster.
12. **C.** A communication plan includes methods used to communicate with response team members, employees, suppliers, and customers. Although not available as a possible answer, it would also include methods used to respond to media requests, including basic templates. None of the other answers are part of a communication plan. Both DRPs and BCPs might include a list of test plans and procedures. Succession planning clarifies who can make decisions during a disaster. A BIA typically includes a list of scenarios with potential loss statements.
13. **B.** Succession planning clarifies who can make decisions during a disaster and can be documented in a chart listing roles and responsibilities along with a chain of command. IT contingency planning focuses on recovery of IT systems. Continuity of operations planning (COOP) identifies methods, such as alternate sites, that an organization can implement after a disaster. Recovery time objective (RTO) identifies the maximum amount of time it should take to restore a system after an outage.
14. **C.** A tabletop exercise is discussion-based and is typically performed in a classroom or conference room setting. Because this is a meeting led by the business continuity plan (BCP) coordinator, it is a tabletop exercise. Functional exercises are hands-on exercises and include simulations and full-blown tests.
15. **D.** The most likely reason for personnel to turn off a server for testing is to test elements of continuity of operations planning (COOP). This helps determine if the organization can continue to operate despite the outage. A business impact analysis (BIA) is performed before creating business continuity plans, not to test them. Succession planning identifies a chain of command during a disaster. Tabletop exercises are discussion-based exercises and do not include manipulating any systems.
16. **D.** Failing humidity controls can cause damage from electrostatic discharge (ESD) if humidity is too low and water damage from condensation if humidity gets too high. Humidity controls do not provide any protection against electromagnetic interference (EMI), temperature, or ventilation.
17. **B.** Increasing the heating, ventilation, and air conditioning (HVAC) capacity results in higher mean time between failures (MTBF) times by keeping systems at lower temperatures. Lower MTBF times indicate more failures. Mean time to recover (MTTR) is unrelated to failures or HVAC systems.
18. **B.** An attacker could manipulate environmental controls to change the temperature or humidity within a data center and cause significant damage. An attacker could block video surveillance by manipulating a video surveillance system, but this wouldn't cause extensive physical damage. Modifying the firewall access control lists (ACLs) or intrusion detection system (IDS) settings might allow remote attacks, but not physical damage.
19. **A.** In this scenario, the most likely reason that the door opened was because the access system was designed to fail-open for personnel safety. If the system was designed to fail-close, then employees would be trapped inside during a fire or other disaster. Nothing in the scenario indicates the system was improperly installed. A fire extinguisher would not include a proximity badge, and it wouldn't work if the proximity reader was destroyed.

20. **A.** Electromagnetic interference (EMI) shielding provides protection against interference from electromagnetic sources such as fluorescent lights. Fencing, video surveillance, and motion detection are all physical security controls.

Chapter 10

Understanding Cryptography

CompTIA Security+ objectives covered in this chapter:

1.4 Given a scenario, implement common protocols and services.

- TLS, SSL, HTTPS

2.9 Given a scenario, select the appropriate control to meet the goals of security.

- Confidentiality (Encryption, Steganography)
- Integrity (Hashing, Digital signatures, Certificates, Non-repudiation)

6.1 Given a scenario, utilize general cryptography concepts.

- Symmetric vs. asymmetric, Session keys, In-band vs. out-of-band key exchange, Fundamental differences and encryption methods (Block vs. stream), Transport encryption, Non-repudiation, Hashing, Key escrow, Steganography, Digital signatures, Use of proven technologies, Elliptic curve and quantum cryptography, Ephemeral key, Perfect forward secrecy

6.2 Given a scenario, use appropriate cryptographic methods.

- MD5, SHA, RIPEMD, AES, DES, 3DES, HMAC, RSA, Diffie-Hellman, RC4, One-time pads, NTLM, NTLMv2, Blowfish, PGP/GPG, TwoFish, DHE, ECDHE, Comparative strengths and performance of algorithms
- Use of algorithms with transport encryption (SSL, TLS, IPSec, SSH, HTTPS)
- Cipher suites (Strong vs. weak ciphers)
- Key stretching (PBKDF2, Bcrypt)

6.3 Given a scenario, use appropriate PKI, certificate management and associated components.

- Certificate authorities and digital certificates (CA, CRLs, OCSP, CSR)
- PKI, Recovery agent, Public key, Private key, Registration, Key escrow, Trust models

Although cryptography is only 12 percent of the exam, you might find that many of these topics aren't as familiar to you as other topics, and you might have to spend more than 12 percent of your study time here. When tackling these topics, don't lose sight of the basics. The first section in this chapter, "Introducing Cryptography Concepts," outlines and summarizes these basics. Other sections dig into the details of hashing, encryption, and Public Key Infrastructure (PKI) components.

Introducing Cryptography Concepts

Cryptography has several important concepts that you need to grasp for the CompTIA Security+ exam, but the topics are often new to many information technology (IT) professionals. Two core topics are integrity and confidentiality (introduced in Chapter 1, “Mastering Security Basics,” as part of the security triad). As an introduction, the following points identify the important core cryptography concepts:

- *Integrity* provides assurances that data has not been modified. Hashing ensures that data has retained integrity.
 - A *hash* is a number derived from performing a calculation on data, such as a message, patch, or update file.
 - Hashing creates a fixed-size string of bits or hexadecimal characters, which cannot be reversed to re-create the original data.
 - Common hashing algorithms include MD5 and Secure Hash Algorithm (SHA).
- *Confidentiality* ensures that data is only viewable by authorized users. Encryption protects the confidentiality of data.
 - *Encryption* scrambles, or ciphers, data to make it unreadable if intercepted. Encryption normally includes an algorithm and a key.
 - *Symmetric encryption* uses the same key to encrypt and decrypt data.
 - *Asymmetric encryption* uses two keys (public and private) created as a matched pair.
 - Anything encrypted with the public key can only be decrypted with the matching private key.
 - Anything encrypted with the private key can only be decrypted with the matching public key.
 - Stream ciphers encrypt data one bit at a time. Block ciphers encrypt data in blocks.
 - Steganography provides a level of confidentiality by hiding data within other files. For example, it's possible to embed data within the white space of a picture file.
- *Authentication* validates an identity.
- *Non-repudiation* prevents a party from denying an action.
- *Digital signatures* provide authentication, non-repudiation, and integrity.
 - Users sign emails with a digital signature, which is a hash of an email message encrypted with the sender's private key.
 - Only the sender's public key can decrypt the hash, providing verification it was encrypted with the sender's private key.

Providing Integrity with Hashing

You can verify integrity with hashing. Hashing is an algorithm performed on data such as a file or message to produce a number called a hash (sometimes called a checksum). The hash is used to verify that data is not modified, tampered with, or corrupted. In other words, you can verify the data has maintained integrity.

A key point about a hash is that no matter how many times you execute the hashing algorithm against the data, the hash will always be the same as long as the data is the same.

Hashes are created at least twice so that they can be compared. For example, imagine a software company is releasing a patch for an application that customers can download. They can calculate the hash of the patch and post both a link to the patch file and the hash on the company site. They might list it as:

- **Patch file.** Patch_v2_3.zip
- **MD5 checksum.** 9d2cf3770edbb49461788164af2331f3

The Message Digest 5 (MD5) checksum is the calculated hash displayed in hexadecimal. Customers can download the hash and then calculate the hash on the downloaded file. If the calculated hash is the same as the hash posted on the web site, it verifies the file has retained integrity. In other words, the file has not changed.

Remember this

Hashing verifies integrity for data such as email, downloaded files, and files stored on a disk. A hash is a number created with a hashing algorithm, and is sometimes listed as a checksum.

MD5

Message Digest 5 (MD5) is a common hashing algorithm that produces a 128-bit hash. Hashes are commonly shown in hexadecimal format instead of a stream of 1s and 0s. For example, the MD5 hash for the patch file (listed as the MD5 checksum) is displayed as 32 hexadecimal characters instead of 128 bits. Hexadecimal characters are composed of four bits and use the numbers 0 through 9 and the characters a through f.

Many applications use MD5 to verify the integrity of files. This includes email, files stored on disks, files downloaded from the Internet, executable files, and more. The “Hashing Files” section shows how you can manually calculate hashes.

SHA

Secure Hash Algorithm (SHA) is another hashing algorithm. There are several variations of SHA grouped into four families: SHA-0, SHA-1, SHA-2, and SHA-3:

- SHA-0 is not used.
- SHA-1 is an updated version that creates 160-bit hashes. This is similar to the MD5 hash except that it creates 160-bit hashes instead of 128-bit hashes.
- SHA-2 improved SHA-1 to overcome potential weaknesses. It includes four versions: SHA-224, SHA-256, SHA-384, and SHA-512. The numbers represent the number of bits in the hash. For example, SHA-256 creates 256-bit hashes.
- SHA-3 uses a different method than SHA-2 and can be used instead. It includes multiple versions with hashes of 224 bits, 256 bits, 384 bits, and 512 bits.

Just as MD5 is used to verify the integrity of files, SHA also verifies file integrity. As an example, it's rare for executable files to be modified. However, some malware modifies executable files by adding malicious code into the file. Rootkits will often modify system-level files.

Some host-based intrusion detection system (HIDS) and antivirus software capture hashes of files on a system when they first scan it and include valid hashes of system files in signature definition files. When they scan a system again, they can capture hashes of executable and system files and compare them with known good hashes. If the hashes are different for an executable or system file, it indicates the file has been modified, and it may have been modified by malware.

HMAC

Another method used to provide integrity is with a Hash-based Message Authentication Code (HMAC). An HMAC is a fixed-length string of bits similar to other hashing algorithms such as MD5 and SHA-1 (known as HMAC-MD5 and HMAC-SHA1). However, HMAC also uses a shared secret key to add some randomness to the result and only the sender and receiver know the secret key.

For example, imagine that one server is sending a message to another server using HMAC-MD5. It starts by first creating a hash of a message with MD5 and then uses a secret key to complete another calculation on the hash. The server then sends the message and the HMAC-MD5 hash to the second server. The second server performs the same calculations and compares the received HMAC-MD5 hash with its result. Just as with any other hash comparison, if the two hashes are the same, the message retained integrity, but if the hashes are different, the message lost integrity.

The HMAC provides both integrity and authenticity of messages. The MD5 portion of the hash provides integrity just as MD5 does. However, because only the server and receiver know the secret key, if the receiver can calculate the same HMAC-MD5 hash as the sender, it knows that the sender used the same key. If an attacker was trying to impersonate the sender, the message wouldn't pass this authenticity check because the attacker wouldn't have the secret key. Internet Protocol security (IPsec) and Transport Layer Security (TLS) often use a version of HMAC such as HMAC-MD5 and HMAC-SHA1.

Remember this

Two popular hashing algorithms used to verify integrity are MD5 and SHA. HMAC verifies both the integrity and authenticity of a message with the use of a shared secret. Other protocols such as IPsec and TLS use HMAC-MD5 and HMAC-SHA1.

Hashing Files

Many applications calculate and compare hashes automatically without any user intervention. For example, digital signatures (described later) use hashes within email, and email applications automatically create and compare the hashes.

Additionally, there are several applications you can use to manually calculate hashes. As an example, *md5sum.exe* is a free program anyone can use to create hashes of files. A Google search on “download md5sum” will show several locations. It runs the MD5 hashing algorithm against a file to create the hash.

Imagine that you downloaded a patch file from a vendor’s site. Before posting the files to the web site, an administrator used a program (such as md5sum) to calculate the hash of the file and posted the hash as an MD5 checksum of *367f0ed4ecd70aefc290d1f7dcb578ab*.

After downloading the file, you can calculate the hash to verify the file hasn’t lost integrity. Figure 10.1 shows one method of calculating the hash of the file. In the figure, I first used the **dir** command to list the two files in the directory (*applicationPatch.exe* and *md5sum.exe*). I then ran **md5sum** against the patch file three times. Each time, md5sum created the same hash *367f0ed4ecd70aefc290d1f7dcb578ab*.



```
C:\SecurityLabs\Md5sum>dir
Volume in drive C is OS
Volume Serial Number is 2E1F-BB37

Directory of C:\SecurityLabs\Md5sum

09/03/2014  12:40 PM    <DIR>          .
09/03/2014  12:40 PM    <DIR>          ..
07/26/2011  08:47 AM             36,531,492 applicationPatch.exe
05/12/2011  07:36 AM             49,152 md5sum.exe
               2 File(s)      36,580,644 bytes
               2 Dir(s)      941,851,095,040 bytes free

C:\SecurityLabs\Md5sum>md5sum applicationPatch.exe
367f0ed4ecd70aefc290d1f7dcb578ab *applicationPatch.exe 1
C:\SecurityLabs\Md5sum>md5sum applicationPatch.exe
367f0ed4ecd70aefc290d1f7dcb578ab *applicationPatch.exe 2
C:\SecurityLabs\Md5sum>md5sum applicationPatch.exe
367f0ed4ecd70aefc290d1f7dcb578ab *applicationPatch.exe 3
C:\SecurityLabs\Md5sum>
```

Figure 10.1: Calculating a hash with md5sum

The figure demonstrates two important points:

- **The hash will always be the same no matter how many times you calculate it.**
In the figure, I ran md5sum three times, but it would give me the same result if I ran it 3,000 times.
- **Hashing verifies the file has retained integrity.**

Because the calculated hash is the same as the MD5 checksum posted on the vendor’s site, it verifies the file has not lost integrity.

In contrast, if md5sum created a different hash than the one posted on the web site, I’d know that the file lost integrity. I wouldn’t necessarily know *why* the file lost integrity. An attacker may have infected it with malware, or it may have lost a bit or two during the transfer. However, I do know that the integrity of the file is lost and the file should not be trusted.

It’s worth stressing that hashes are one-way functions. In other words, you can calculate a hash on a file or a message, but you can’t use the hash to reproduce the original data. The hashing algorithms always create a fixed-size bit string regardless of the size of the original data.

As an example, the MD5 hash from the message “I will pass the Security+ exam” is: *5384128261CF2EEA6D90ADACE48CD41B*. However, you can’t look at the hash and identify the message, or even know that it is a hash of a message.

The hash shown in Figure 10.1 was calculated on a 35 MB executable file. However, the hash doesn’t give you a clue about the size of the file, the type of the file, or anything else. It could just as easily be a single sentence message, a 10 KB email, a 7 GB database file, or something else.

If you want to work with hashes yourself, check out the Hashing Lab in the online exercises for this book at <http://gcgapremium.com/labs/>.

Hashing Passwords

Passwords are often stored as hashes. When a user creates a new password, the system calculates the hash for the password and then stores the hash. Later, when the user authenticates by entering a username and password, the system calculates the hash of the entered password, and then compares it with the stored hash. If the hashes are the same, it indicates that the user entered the correct password.

As mentioned in Chapter 7, “Identifying Advanced Attacks,” many password attacks attempt to discover a password by calculating a hash on a guessed password, and then comparing it with the stored hash of the password. Many systems had additional characters as salts to passwords to thwart these attacks.

Remember this

Hashing is a one-way function that creates a string of characters. You cannot reverse the hash to re-create the original file. Passwords are often stored as hashes instead of storing the actual password. Additionally, applications often salt passwords with extra characters before hashing them.

Hashing Messages

Hashing provides integrity for messages. It provides assurance to someone receiving a message that the message has not been modified. Imagine that Lisa is sending a message to Bart, as shown in Figure 10.2. The message is “The price is \$75.” This message is not secret, so there is no need to encrypt it. However, we do want to provide integrity, so this explanation is focused only on hashing.

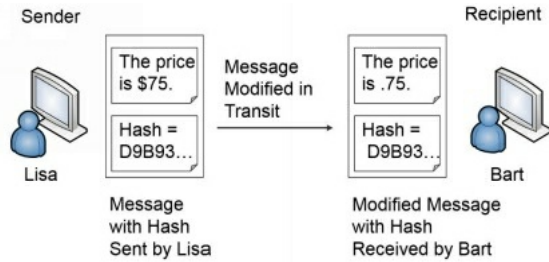


Figure 10.2: Simplified hash process

An application on Lisa’s computer calculates the MD5 hash as *D9B93C99B62646ABD06C887039053F56*. In the figure, I’ve shortened the full hash down to just the first five characters of “*D9B93*.” Lisa then sends both the message and the hash to Bart.

In this example, something modified the message before it reaches Bart. When Bart receives the message and the original hash, the message is now “The price is .75.” Note that the message is modified in transit, but the hash is *not* modified.

A program on Bart’s computer calculates the MD5 hash on the received message as *564294439E1617F5628A3E3EB75643FE*. It then compares the received hash with the calculated hash:

- Hash created on Lisa’s computer, and received by Bart’s computer:
D9B93C99B62646ABD06C887039053F56
- Hash created on Bart’s computer:
564294439E1617F5628A3E3EB75643FE

Clearly, the hashes are different, so you know the message lost integrity. The program on Bart’s computer would report the discrepancy. Bart doesn’t know what caused the problem. It could have been a malicious attacker changing the message, or it could have been a technical problem. However, Bart does know the received message isn’t the same as the sent message and he shouldn’t trust it.

Using HMAC

You might have noticed a problem in the explanation of the hashed message. If an attacker can change the message, why can't the attacker change the hash, too? In other words, if hacker Harry changed the message to "The price is .75," he could also calculate the hash on the modified message and replace the original hash with the modified hash. Here's the result:

- Hash created on Lisa's computer:
D9B93C99B62646ABD06C887039053F56
- Modified hash inserted by attacker after modifying the message:
564294439E1617F5628A3E3EB75643FE
- Hash created for modified message on Bart's computer:
564294439E1617F5628A3E3EB75643FE

The calculated hash on the modified message would be the same as the received hash. This erroneously indicates that the message maintained integrity. HMAC helps solve this problem.

With HMAC, both Lisa and Bart's computers would know the same secret key and use it to create an HMAC-MD5 hash instead of just an MD5 hash. Figure 10.3 shows the result.

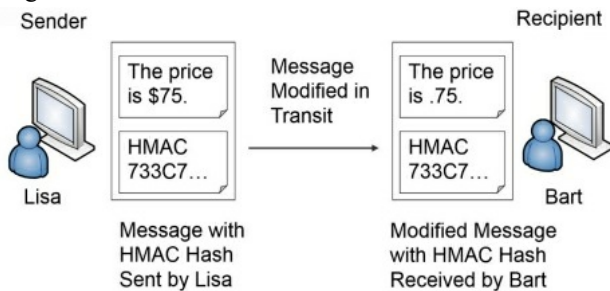


Figure 10.3: Using HMAC

Lisa is still sending the same message. The MD5 hash is *D9B93C99B62646ABD06C887039053F56*. However, after applying the HMAC secret key, the HMAC-MD5 hash is *733C70A54A13744D5C2C9C4BA3B15034*. For brevity, I shortened this to only the first five characters (733C7) in the figure.

An attacker can modify the message in transit just as before. However, the attacker doesn't know the secret key, so he can't calculate the HMAC hash.

Bart's computer calculates the HMAC-MD5 hash on the received message using the shared secret key. It then compares the calculated hash with the hash received from Lisa:

- HMAC-MD5 hash created on Lisa's computer:
733C70A54A13744D5C2C9C4BA3B15034
- HMAC-MD5 hash created on Bart's computer:
1B4FF0F6C04434BF97F1E3DDD4B6C137

Again, you can see that the hashes are different and the message has lost integrity. If the messages weren't modified, the HMAC-MD5 hash would be the same.

Table 10.1 summarizes the important hashing protocols covered in the CompTIA Security+ exam.

Algorithm	Type	Comments
MD5	Hashing - Integrity	Creates 128-bit hashes
SHA-1	Hashing - Integrity	Creates 160-bit hashes
SHA-2	Hashing - Integrity	Creates 224-, 256-, 384-, or 512-bit hashes
HMAC-MD5	Integrity/Authenticity	Creates 128-bit hashes
HMAC-SHA1	Integrity/Authenticity	Creates 160-bit hashes

Table 10.1: Hashing protocols

Remember this

If you can recognize the hashing algorithms such as MD5, SHA, and HMAC, it will help you answer many exam questions. For example, if a question asks what you would use to encrypt data and lists hashing algorithms, you can quickly eliminate them because they don't hash data.

Other Hash Algorithms

MD5 and SHA are the most popular hashing algorithms in use today. However, the CompTIA Security+ objectives mention other hashing algorithms that aren't used as often. The following sections describe these other hashing algorithms.

RIPEMD

RACE Integrity Primitives Evaluation Message Digest (RIPEMD) is another hash function. Different versions create different size hashes. RIPEMD-160 creates 160-bit, fixed-size hashes. Other versions create hash sizes of 128 bits, 256 bits, and 320 bits.

LANMAN and NTLM

Older Microsoft systems used hashing algorithms to secure passwords. Although the protocols were primarily used for authentication, they are relevant here due to how they used hashing.

LANMAN

LAN Manager (LANMAN) is a very old authentication protocol used to provide backward compatibility to Windows 95, 98, and ME clients. LANMAN has significant weaknesses with how it stores the password.

Even though LANMAN is very old, some legacy services and software applications still use LANMAN. Additionally, some newer authentication protocols will still store passwords in the LANMAN format if the password is less than 15 characters long. LANMAN cannot handle passwords of 15 characters or more.

LANMAN performs a hashing algorithm on passwords that makes it easy for password-cracking tools, such as L0phtCrack, to discover the actual password. The LANMAN passwords are always stored as 14 characters. If the password is less than 14 characters, it pads the password with trailing spaces. It then converts all lowercase characters to uppercase and creates a hash on each of 7-character strings. The two hashes are stored locally as a single string.

If the password is only 7 characters long, the resulting hash on the trailing seven spaces would always be AAD3B435B51404EE. If the entire password is blank, it becomes two sets of seven spaces. The hash of the entire password is then two sets of AAD3B435B51404EE, or AAD3B435B51404EEAAD3B435B51404EE.

However, even if the password is more than 7 characters long, it doesn't take L0phtCrack long to successfully crack the password because it only needs to work on 7 characters at a time. L0phtCrack guesses different 1- to 7-character combinations and compares the hash of the guessed password with the hash of the stored password.

For the best security, LANMAN should be disabled on all computers within a network. The following operating systems ship with LANMAN disabled by default: Windows Vista, Windows 7, Windows 8, Windows Server 2008, Server 2008 R2, Server 2012, and Server 2012 R2.

NTLM

Microsoft introduced NT LAN Manager (NTLM) as an improvement over LANMAN. There are two versions of NTLM: NTLM (or NTLMv1) and NTLMv2.

NTLMv1 uses an MD4 hash of the user's password, and for backward compatibility it also uses the LANMAN hash if the password is 14 characters or less. Both LANMAN and MD4 are considered compromised, resulting in known vulnerabilities with NTLMv1 today.

Microsoft improved NTLMv1 with NTLMv2. NTLMv2 uses a completely different process and uses the more secure MD5 algorithm. NTLMv2 is significantly complex, making it infeasible to crack using current technologies.

Although NTLMv1 and NTLMv2 provide improvements over LANMAN, a significant vulnerability exists in systems before Windows Vista such as Windows XP and Windows Server 2003. Specifically, LANMAN is still enabled by default on older systems. When it is enabled, these systems use the LANMAN hash for backward compatibility, in addition to the more secure NTLMv2.

Many networks impose policies requiring administrator passwords to be 15 characters or more. This 15-character requirement overcomes the LANMAN vulnerability because LANMAN cannot create hashes of passwords with more than 14 characters. If an organization doesn't need the LANMAN passwords for backward compatibility, it can disable it. This article provides details on how to do so:

<http://support.microsoft.com/kb/299656>.

Providing Confidentiality with Encryption

Encryption provides confidentiality and prevents unauthorized disclosure of data. Encrypted data is in a ciphertext format that is unreadable. Attackers can't read encrypted traffic sent over a network or encrypted data stored on a system. In contrast, if data is sent in cleartext, an attacker can capture and read the data using a protocol analyzer.

Similarly, you can protect data at rest by encrypting it. Chapter 5, "Securing Hosts and Data," introduced several levels of encryption. For example, if a customer database includes credit card data, you can encrypt the fields holding this data, but you don't necessarily have to encrypt the entire database file. It's also possible to encrypt files and entire disks.

The two primary encryption methods are symmetric and asymmetric. Symmetric encryption encrypts and decrypts data with the same key. Asymmetric encryption encrypts and decrypts data using a matched key pair of a public key and a private key.

These encryption methods include two elements:

- **Algorithm.** The algorithm performs mathematical calculations on data. The algorithm is always the same.
- **Key.** The key is a number that provides variability for the encryption. It is either kept private and/or changed frequently.

Remember this

Encryption provides confidentiality and helps ensure that data is viewable only by authorized users. This applies to any type of data, including data at rest, such as data stored in a database, or data in transit sent over a network.

Symmetric Encryption

Symmetric encryption uses the same key to encrypt and decrypt data. In other words, if you encrypt data with a key of three, you decrypt it with the same key of three. Symmetric encryption is also called secret-key encryption or session-key encryption.

As a simple example, when I was a child, a friend and I used to pass encoded messages back and forth to each other. Our algorithm was:

- **Encryption algorithm.** Move X spaces forward to encrypt.
- **Decryption algorithm.** Move X spaces backward to decrypt.

On the way to school, we would identify the key (X) we would use that day. For example, we may have used the key of three one day. If I wanted to encrypt a message, I would move each character three spaces forward, and he would decrypt the message by moving three spaces backward.

Imagine the message “PASS” needs to be sent.

- Three characters past “P” is “S”—Start at P (Q, R, S)
- Three characters past “A” is “D”—Start at A (B, C, D)
- Three characters past “S” is “V”—Start at S (T, U, V)
- Three characters past “S” is “V”—Start at S (T, U, V)

The encrypted message is SDVV. My friend decrypted it by moving backward three spaces and learning that “PASS” is the original message.

This shows how symmetric encryption uses the same key for encryption and decryption. If I encrypted the message with a key of three, my friend wouldn’t be able to decrypt it with anything but a key of three. It also helps to demonstrate an algorithm and a key, though it is admittedly simple. Most algorithms and keys are much more complex. For example, the Advanced Encryption Standard (AES) symmetric algorithm typically uses 128-bit keys, but can use keys with 256 bits.

Sophisticated symmetric encryption techniques use the same components of an algorithm and a key. Imagine two servers sending encrypted traffic back and forth to each other using AES symmetric encryption. They both use the same AES algorithm and the same key for this data. The data is encrypted on one server with AES and a key, sent over the wire or other transmission medium, and the same key is used to decrypt it on the other server. Similarly, if a database includes encrypted data, the key used to encrypt data is the same key used to decrypt data.

However, symmetric encryption doesn’t use the same key to encrypt and decrypt all data. For example, my friend and I used a different key each day. On the way to school, we decided on a key to use for that day. The next day, we picked a different key. If someone cracked our code yesterday, they couldn’t easily crack our code today.

Symmetric encryption algorithms change keys much more often than once a day. For example, imagine an algorithm uses a key of 123 to encrypt a project file. It could then use a key of 456 to encrypt a spreadsheet file. The key of 123 can only decrypt the project file and the key of 456 can only decrypt the spreadsheet file.

On the other hand, if symmetric encryption always used the same key of 123, it would add vulnerabilities. First, when keys are reused, the encryption is easier to crack. Second, once the key is cracked, all data encrypted with this key is compromised. If attackers discover the key of 123, not only would they have access to the project file, but they would also have access to the spreadsheet file and any other data encrypted with this same key.

As a more realistic example, Chapter 1 describes how Remote Authentication Dial-In User Service (RADIUS) encrypts password packets. RADIUS uses shared keys for symmetric encryption. When users authenticate, RADIUS servers and clients use the shared key to encrypt and decrypt a message such as a pseudorandom number in a challenge/response session. Without the shared key, clients are unable to decrypt the message and respond appropriately.

Remember this

Symmetric encryption uses the same key to encrypt and decrypt data. For example, when transmitting encrypted data, symmetric encryption algorithms use the same key to encrypt and decrypt data at both ends of the transmission media. RADIUS uses symmetric encryption.

Block Versus Stream Ciphers

Most symmetric algorithms use either a block cipher or a stream cipher. They are both symmetric, so they both use the same key to encrypt or decrypt data. However, they divide data in different ways.

A block cipher encrypts data in specific-sized blocks, such as 64-bit blocks or 128-bit blocks. The block cipher divides large files or messages into these blocks and then encrypts each individual block separately. Stream ciphers encrypt data as a stream of bits or bytes rather than dividing it into blocks.

In general, stream ciphers are more efficient than block ciphers when the size of the data is unknown or sent in a continuous stream, such as when streaming audio and video over a network. Block ciphers are more efficient when the size of the data is known, such as when encrypting a file or a specific-sized database field.

An important principle when using a stream cipher is that encryption keys should never be reused. If a key is reused, it is easier to crack the

encryption.

For example, Chapter 4, “Securing Your Network,” introduced Wired Equivalent Privacy (WEP) and initialization vector (IV) attacks. WEP uses Rivest Cipher 4 (RC4) stream cipher for symmetric encryption. RC4 is a secure algorithm when it’s implemented correctly, but WEP did not follow the important stream cipher principle of never reusing keys. If wireless systems generate enough traffic, WEP reuses keys for RC4. Attackers discovered they could use packet injection techniques to increase the number of packets on a wireless network, detect the duplicate keys, and crack the encryption.

Remember this

Stream ciphers encrypt data a single bit, or a single byte, at a time in a stream. Block ciphers encrypt data in a specific-sized block such as 64-bit or 128-bit blocks. Stream ciphers are more efficient than block ciphers when encrypting data in a continuous stream.

Comparing Symmetric Encryption to a Door Key

Occasionally, security professionals compare symmetric keys to a house key, and this analogy helps some people understand symmetric encryption a little better. For example, imagine Marge moves into a new home. She'll receive a single key that she can use to lock and unlock her home. Of course, Marge can't use this key to unlock her neighbor's home.

Later, Marge marries Homer, and Homer moves into Marge's home. Marge can create a copy of her house key and give it to Homer. Homer can now use that copy of the key to lock and unlock the house. By sharing copies of the same key, it doesn't matter whether Marge or Homer is the one who locks the door; they can both unlock it.

Similarly, symmetric encryption uses a single key to encrypt and decrypt data. If a copy of the symmetric key is shared, others who have the key can also encrypt and decrypt data.

. . .

AES

The Advanced Encryption Standard (AES) is a strong symmetric block cipher that encrypts data in 128-bit blocks. The National Institute of Standards and Technology (NIST) adopted AES from the Rijndael encryption algorithm after a lengthy evaluation of several different algorithms. NIST is a U.S. agency that develops and promotes standards. They spent about five years conducting a review of 15 different symmetric algorithms and identified AES as the best of the 15.

AES can use key sizes of 128 bits, 192 bits, or 256 bits, and it's sometimes referred to as AES-128, AES-192, or AES-256 to identify how many bits are used in the key. When more bits are used, it makes it more difficult to discover the key. AES-128 provides strong protection, but AES-256 provides stronger protection.

Because of its strengths, AES has been adopted in a wide assortment of applications. For example, many applications that encrypt data on USB drives use AES. Some of the strengths of AES are:

- **Fast.** AES uses elegant mathematical formulas and only requires one pass to encrypt and decrypt data. In contrast, 3DES requires multiple passes to encrypt and decrypt data.
- **Efficient.** AES is less resource intensive than other encryption algorithms such as 3DES. AES encrypts and decrypts quickly even when ciphering data on small devices, such as USB flash drives.
- **Strong.** AES provides strong encryption of data providing a high level of confidentiality.

Remember this

AES is a strong symmetric block cipher that encrypts data in 128-bit blocks. AES uses 128-bit, 192-bit, or 256-bit keys.

DES

Data Encryption Standard (DES) is a symmetric block cipher that was widely used for many years, dating back to the 1970s. It encrypts data in 64-bit blocks. However, it uses a relatively small key of only 56 bits and can be broken with brute force attacks. In the '70s, the technology required to break 56-bit encryption wasn't easily available, but with the advances in computer technology, a 56-bit key is now considered trivial. DES is not recommended for use today.

3DES

3DES (pronounced as "Triple DES") is a symmetric block cipher designed as an improvement over the known weaknesses of DES. In basic terms, it encrypts data using the DES algorithm in three separate passes and uses multiple keys. Just as DES encrypts data in 64-bit blocks, 3DES also encrypts data in 64-bit blocks.

Although 3DES is a strong algorithm, it isn't used as often as AES today. AES is much less resource intensive. However, if hardware doesn't support AES, 3DES is a suitable alternative. 3DES uses key sizes of 56 bits, 112 bits, or 168 bits.

Remember this

DES and 3DES are block ciphers that encrypt data in 64-bit blocks. 3DES was originally designed as a replacement for DES, but NIST selected AES as the current standard. However, 3DES is still used in some applications, such as when legacy hardware doesn't support AES.

RC4

Ron Rivest invented several versions of RC, which are sometimes referred to as Ron's Code or Rivest Cipher. The most commonly used version is RC4 (also called ARC4), which is a symmetric stream cipher and it can use between 40 and 2,048 bits.

It's worthwhile pointing out that this is the same RC4 used in WEP. WEP's vulnerabilities weren't because it used RC4, but instead because it did not follow a basic rule of a stream cipher: Don't reuse keys.

When implemented correctly, RC4 has enjoyed a long life as a strong cipher. For many years, it has been the recommended encryption mechanism in Secure Sockets Layer (SSL) and Transport Layer Security (TLS). SSL and TLS encrypt Hypertext Transfer Protocol Secure (HTTPS) connections on the Internet.

However, experts have speculated since 2013 that agencies such as the U.S. National Security Agency (NSA) can break RC4, even when implemented correctly such as in TLS. Because of this, companies such as Microsoft recommend disabling RC4 and using AES instead. Even though AES is a block cipher and RC4 is a stream cipher, TLS can implement either one.

Blowfish and Twofish

Blowfish is a strong symmetric block cipher that is still widely used today. It encrypts data in 64-bit blocks and supports key sizes between 32 and 448 bits. Bruce Schneier (a widely respected voice in IT security) designed Blowfish as a general-purpose algorithm to replace DES.

Interestingly, Blowfish is actually faster than AES in some instances. This is especially true when comparing Blowfish with AES-256. Part of the reason is that Blowfish encrypts data in smaller 64-bit blocks, whereas AES encrypts data in 128-bit blocks.

Twofish is related to Blowfish, but it encrypts data in 128-bit blocks and it supports 128-, 192-, or 256-bit keys. It was one of the finalist algorithms evaluated by NIST for AES. However, NIST selected another algorithm (Rijndael) as AES.

Remember this

RC4 is a strong symmetric stream cipher. Blowfish is a 64-bit block cipher and Twofish is a 128-bit block cipher. Although NIST recommends AES as the standard, Blowfish is faster than AES-256.

One-Time Pad

The one-time pad cipher has been around since 1917, and many consider it to be one of the most secure algorithms, though it is labor intensive. The one-time pad is a hard copy printout of keys in a pad of paper. Each piece of paper in the pad has a single key along with a serial number that identifies the page. Spies have used these with miniature pads of paper so small it requires a magnifying glass to read the keys.

As an example, imagine Lisa and Bart are two spies and they want to be able to share secrets with each other. They both have identical pads. Lisa can create a message and encrypt it with a key from one of the pages in her pad. She includes the serial number of the page and sends the encrypted message and the serial number to Bart. She then destroys the page she used to create the message.

Bart receives the message with the serial number and he locates the page in his one-time pad. He uses the key on the page to decrypt the message, and after decrypting the message, Bart destroys the key.

One-time pads follow some basic principles for security. The keys must be random, and the key must be the same length as the plaintext message. Users must destroy the key after using it once and never reuse a key. Additionally, there should only be two copies of a one-time pad.

One-time pads have been adapted into some computer applications. For example, a token or fob (described in Chapter 1) is in the *something you have* factor of authentication and makes use of one-time use, rolling passwords similar to one-time pads. As a reminder, users have a token that displays a number on an LCD display. This number changes regularly, such as every 60 seconds. The fob is synchronized with a server that will know the displayed number at any point in time.

Table 10.2 summarizes the important symmetric protocols covered in the CompTIA Security+ exam.

Algorithm	Type	Method	Key Size
AES	Symmetric encryption	128-bit block cipher	128-, 192-, or 256-bit key
DES	Symmetric encryption	64-bit block cipher	56-bit key
3DES	Symmetric encryption	64-bit block cipher	56-, 112-, or 168-bit key
Blowfish	Symmetric encryption	64-bit block cipher	32- to 448-bit key
Twofish	Symmetric encryption	128-bit block cipher	128-, 192-, or 256-bit key
RC4	Symmetric encryption	Stream cipher	40- to 2,048-bit key

Table 10.2: Symmetric encryption protocols

Remember this

If you can recognize the symmetric algorithms such as AES, DES, 3DES, Blowfish, and Twofish, it will help you answer many exam questions. For example, if a question asks what you would use to hash and it lists encryption algorithms, you can quickly eliminate them because they don't encrypt data. You should also

know the size of the blocks and the size of the keys listed in Table 10.2.

Asymmetric Encryption

Asymmetric encryption uses two keys in a matched pair to encrypt and decrypt data—a public key and a private key. There are several important points to remember with these keys:

- If the *public* key encrypts information, only the matching *private* key can decrypt the same information.
- If the *private* key encrypts information, only the matching *public* key can decrypt the same information.
- Private keys are always kept private and never shared.
- Public keys are freely shared by embedding them in a certificate.

Remember this

Only a private key can decrypt information encrypted with a matching public key. Only a public key can decrypt information encrypted with a matching private key. A key element of several asymmetric encryption methods is that they require a certificate and a PKI.

Although asymmetric encryption is very strong, it is also very resource intensive. It takes a significant amount of processing power to encrypt and decrypt data, especially when compared with symmetric encryption. Most cryptographic protocols that use asymmetric encryption only use it to privately share a symmetric key. They then use symmetric encryption to encrypt and decrypt data because symmetric encryption is so much more efficient.

Some of the more advanced topics related to asymmetric encryption become harder to understand if you don't understand the relationship of matched public and private key pairs. However, because you can't actually see these keys, the concepts are hard to grasp for some people. The Rayburn box demonstrates how you can use physical keys for the same purposes as these public and private keys.

The Rayburn Box

I often talk about the Rayburn box in the classroom to help people understand the usage of public and private keys. A Rayburn box is a lockbox that allows people to securely transfer items over long distances. It has two keys. One key can lock the box, but can't unlock it. The other key can unlock the box, but can't lock it.

Both keys are matched to one box and won't work with other boxes:

- Only one copy of one key exists—think of it as the private key.
- Multiple copies of the other key exist, and copies are freely made and distributed—think of these as public keys.

The box comes in two different versions. In one version, it's used to send secrets in a confidential manner to prevent unauthorized disclosure. In the other version, it's used to send messages with authentication, so you know the sender actually sent the message and that the message wasn't modified in transit.

The Rayburn Box Used to Send Secrets

Imagine that I wanted you to send some proprietary information and a working model of a new invention to me. Obviously, we wouldn't want anyone else to be able to access the information or the working model. I could send you the empty open box with a copy of the key used to lock it.

You place everything in the box and then lock it with the public key I've sent with the box. This key can't unlock the box, so even if other people had copies of the public key that I sent to you, they couldn't use it to unlock the box. When I receive the box from you, I can unlock it with the only key that will unlock it—my private key.

This is similar to how public and private keys are used to send encrypted data over the Internet to ensure confidentiality. The public key encrypts information. Information encrypted with a public key can only be decrypted with the matching private key. Many copies of the public key are available, but only one private key exists, and the private key always stays private. The “Encrypting HTTPS Traffic with SSL or TLS” section later in this chapter shows this process in more depth.

The Rayburn Box Used for Authentication

With a little rekeying of the box, I can use it to send messages while giving assurances to recipients that I sent the message. In this context, the message isn't secret and doesn't need to be protected. Instead, it's important that you know I sent the message.

When used this way, the private key will lock the Rayburn box, but it cannot unlock the box. Instead, only a matching public key can unlock it. Multiple copies of the public key exist and anyone with a public key can unlock the box. However, after unlocking the box with a matching public key, it isn't possible to lock it with the public key.

Imagine that you and I are allies in a battle. I want to give you a message of “SY0-401,” which is a code telling you to launch a specific attack at a specific time. We don't care if someone reads this message because it's a code. However, we need you to have assurances that I sent the message.

I write the message, place it in the box, and lock it with my private key. When you receive it, you can unlock it with the matching public key. Because the public key opens it, you know this is my box and it was locked with my private key—you know I sent the message.

If someone intercepted the box and opened it with the public key, he or she wouldn't be able to lock it again using the public key, so you'd receive an open box. An open box with a message inside it doesn't prove I sent it. The only way you know that I sent it is if you receive a locked box that you can unlock with the matching public key.

This is similar to how digital signatures use public and private keys. The “Signing Email with Digital Signatures” section later in this chapter explains digital signatures in more depth. In short, I can send you a message digitally signed with my private key. If you can decrypt the digital signature with my matching public key, you know it was encrypted, or signed, with my private key. Because only one copy of the private key exists, and I'm the only person who can access it, you know I sent the message.

The Rayburn Box Demystified

Before you try to find a Rayburn box, let me clear something up. The Rayburn box is just a figment of my imagination. Rayburn is my middle name.

I haven't discovered a real-world example of how public/private keys work, so I've created the Rayburn box as a metaphor to help people visualize how public/private keys work. Feel free to build one if you want.

Certificates

A key element of asymmetric encryption is a certificate. A certificate is a digital document that includes the public key and information on the owner of the certificate. Certificate authorities (CAs) issue and manage certificates. CAs are explored in greater depth later in this chapter. Certificates are used for a variety of purposes, including encryption, authentication, and digital signatures.

Figure 10.4 shows a sample certificate with the public key selected. Users and applications share the certificate file to share the public key. They do not share the private key.

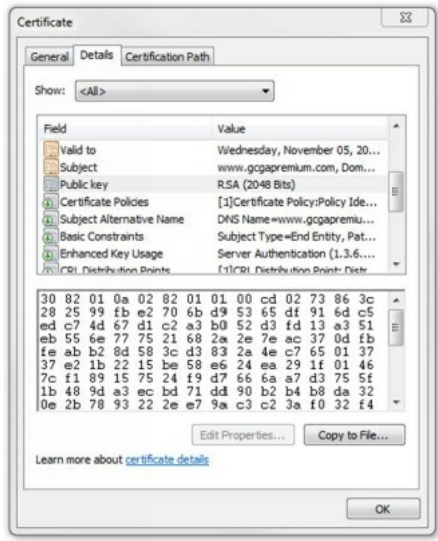


Figure 10.4: Certificate with public key selected

Notice that there is much more information in the certificate than just the public key. Some of it is visible in the figure, but there is more. Common elements within a certificate include:

- **Serial number.** The serial number uniquely identifies the certificate. The CA uses this serial number to validate a certificate. If the CA revokes the certificate, it publishes this serial number in a certificate revocation list (CRL).
- **Issuer.** This identifies the CA that issued the certificate.
- **Validity dates.** Certificates include “Valid From” and “Valid To” dates. This ensures a certificate expires at some point.
- **Subject.** This identifies the owner of the certificate. In the figure, it identifies the subject as the web site with the *www.gcgapremium.com* domain name.
- **Public key.** RSA asymmetric encryption uses the public key in combination with the matching private key.
- **Usage.** Some certificates are only for encryption or authentication, whereas other certificates support multiple usages.

Remember this

Certificates are an important part of asymmetric encryption. Certificates include public keys along with details on the owner of the certificate and on the CA that issued the certificate. Certificate owners share their public key by sharing a copy of their certificate.

If you want to view a certificate, check out the View a Certificate Lab in the online exercises for this book at <http://gcppremium.com/labs/>.

RSA

Ron Rivest, Adi Shamir, and Leonard Adleman developed RSA in 1977 and the acronym uses their last names. It is an asymmetric encryption method using both a public key and a private key in a matched pair, and it is widely used on the Internet and elsewhere due to its strong security.

As an example, email applications often use RSA to privately share a symmetric key between two systems. The application uses the recipient's public key to encrypt a symmetric key, and the recipient's private key decrypts it. The "Protecting Email" section later in this chapter discusses this process in more detail.

Chapter 5 introduces Trusted Platform Modules (TPMs) and hardware security modules (HSMs) used for hardware encryption. As a reminder, TPMs and HSMs provide secure storage for RSA keys.

RSA uses the mathematical properties of prime numbers to generate secure public and private keys. Specifically, RSA relies on the fact that it is difficult to factor the product of two large prime numbers. The math is complex and intriguing to mathematicians, but you don't have to understand the math to understand that RSA is secure.

Researchers published a paper in 2010 identifying how long it took to factor a 232-digit number (768 bits). They wrote that it took them about two and a half years using hundreds of systems. They estimated that if someone used a single 2.2 GHz computer, it would take 1,500 years to complete.

Although the processing power of computers has advanced since 2010, so has the number of bits in keys used by RSA. As of 2014, RSA uses a minimum of 1,024-bit keys. RSA Security (a company that frequently tests the security of RSA) recommends using key sizes of at least 2,048 bits long, and 4,096-bit keys are in use.

Remember this

RSA is widely used to protect data such as email and other data transmitted over the Internet. It uses both a public key and a private key in a matched pair.

Static Versus Ephemeral Keys

The two primary categories of asymmetric keys are static and ephemeral. In general, *static keys* are semipermanent and stay the same over a long period of time. In contrast, *ephemeral keys* have very short lifetimes and are re-created for each session.

RSA uses static keys. A certificate includes an embedded public key matched to a private key and this key pair is valid for the lifetime of a certificate, such as a year. Certificates have expiration dates and systems continue to use these keys until the certificate expires. A benefit of static keys is that a CA can validate them as discussed in the “Validating Certificates” section later in this chapter.

An ephemeral key pair includes a private ephemeral key and a public ephemeral key. However, systems use these key pairs for a single session and then discard them. Some versions of Diffie-Hellman use static keys and some versions use ephemeral keys.

Perfect forward secrecy is an important characteristic that ephemeral keys comply with in asymmetric encryption. Perfect forward secrecy indicates that a cryptographic system generates random public keys for each session and it doesn’t use a deterministic algorithm to do so. In other words, given the same input, the algorithm will create a different public key. This helps ensure that systems do not reuse keys.

Elliptic Curve Cryptography

Elliptic curve cryptography (ECC) is commonly used with small wireless devices because it doesn't take much processing power to achieve the desired security. It uses mathematical equations to formulate an elliptical curve. It then graphs points on the curve to create keys. This is mathematically easier and requires less processing power, while also being more difficult to crack.

The math behind ECC is quite complex, but a simple fact helps to illustrate its strength. In 2005, the U.S. National Security Agency (NSA) announced approval of ECC for digital signatures and Diffie-Hellman key agreements. If the NSA has endorsed and approved ECC, you can bet that it is well tested and strong.

Diffie-Hellman

Diffie-Hellman is a key exchange algorithm used to privately share a symmetric key between two parties. Once the two parties know the symmetric key, they use symmetric encryption to encrypt the data.

Whitfield Diffie and Martin Hellman first published the Diffie-Hellman scheme in 1976. Interestingly, Malcolm J. Williamson secretly created a similar algorithm while working in a British intelligence agency. It is widely believed that the work of these three provided the basis for public-key cryptography.

Diffie-Hellman methods support both static keys and ephemeral keys. RSA is based on the Diffie-Hellman key exchange concepts using static keys. Two Diffie-Hellman methods that use ephemeral keys are:

- **DHE.** Diffie-Hellman Ephemeral (DHE) uses ephemeral keys, generating different keys for each session. Some documents list this as Ephemeral Diffie-Hellman (EDH).
- **ECDHE.** Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) uses ephemeral keys generated using ECC. Another version, Elliptic Curve Diffie-Hellman (ECDH), uses static keys.

Remember this

Diffie-Hellman is a secure method of sharing symmetric encryption keys over a public network. Elliptic curve cryptography is commonly used with small wireless devices. ECDHE is a version of Diffie-Hellman that uses elliptic curve cryptography to generate encryption keys.

Steganography

Steganography hides data inside other data, or, as some people have said, it hides data in plain sight. The goal is to hide the data in such a way that no one suspects there is a hidden message. It doesn't actually encrypt the data, so it can't be classified as either symmetric or asymmetric. However, it can effectively hide information, so it is included with encryption topics.

Imagine if a terrorist in one country wanted to communicate with terrorist cells in another country. They could use steganography by posting graphics to a web page. For example, the web page may include a Graphics Interchange Format (GIF) or Joint Photographic Experts Group (JPEG) picture of a tree. When they want to send a message, they modify the graphic to include the message and post the modified graphic to the site. This looks like the same tree. However, if their contacts know to look for a message, they can download the graphics file and retrieve the message.

Some common examples of steganography are:

- **Hide data by manipulating bits.** It's possible to manipulate some bits within an image or sound file to embed a message. One method of embedding data in large files is modifying the least significant bit in some bytes. By modifying the least significant bit in some of the individual bytes of a JPEG file, it embeds a message, but the changes are so small that they are difficult to detect. However, if people know the file includes a message, they can easily retrieve it.
- **Hide data in the white space of a file.** Many files have unused space (called white space) at the end of file clusters. Imagine a small 6 KB file stored in two 4 KB clusters. It has an extra 2 KB of unused space and it's possible to fill this white space with a message. For example, you can embed a message into the white space of a GIF or JPEG file without altering the file size.

Security professionals use steganalysis techniques to detect steganography, and the most common method is with hashing. If a single bit of a file is modified, the hashing algorithm creates a different hash. By regularly taking the hashes of different files and comparing them with previous hashes, it's easy to detect when a file has been modified.

If you want to see how to embed a text file in an image file, check out the Steganography Lab mentioned in Chapter 1 at

<http://gcgapremium.com/labs/>.

Remember this

Steganography hides messages or other data within a file. For example, you can hide messages within the white space of a JPEG or GIF file. Security professionals use hashing to detect changes in files that may indicate the use of steganography.

Quantum Cryptography

Quantum cryptography is based on quantum physics and photons, the smallest measure of light. Some applications use quantum cryptography to transmit an encryption key as a series of photons.

Normally, a photon spins and emits light in all directions, similar to how a light bulb emits light in all directions. In quantum cryptography, photons are modified so that they spin and emit light in a single direction, such as up and down or side to side. It's then possible to assign a value for the direction.

For example, if the photon emits light up and down, it's a 1. If it emits light side to side, it's a 0. It's then possible to convert a stream of photons into an encryption key. Two parties exchange photon streams identifying and validating the key.

There's much more involved in this photon conversation, but I'm going to skip the quantum physics. However, there is one more important element. When a photon is read or measured, it changes direction. If a third party reads any of the photons in the stream, it will be obvious to the two parties trying to exchange the key. When the two parties realize someone is eavesdropping, they won't exchange secure data until the third party is removed.

Using Cryptographic Protocols

With a basic understanding of hashing, symmetric encryption, and asymmetric encryption, it's easier to grasp how cryptography is used. Many applications use a combination of these methods, and it's important to understand how they're intertwined.

When describing public and private keys earlier, it was stressed that one key encrypts and the other key decrypts. A common question is “which one encrypts and which one decrypts?” The answer depends on what you're trying to accomplish. The following sections describe the details, but as an overview, these are the important points related to these keys:

- Email digital signatures
 - The *sender's private key* encrypts (or signs).
 - The *sender's public key* decrypts.
- Email encryption
 - The *recipient's public key* encrypts.
 - The *recipient's private key* decrypts.
- Web site encryption
 - The *web site's public key* encrypts (a symmetric key).
 - The *web site's private key* decrypts (a symmetric key).
 - The *symmetric key* encrypts data in the web site session.

Email and web site encryption commonly use a combination of both asymmetric and symmetric encryption. They use asymmetric encryption to privately share a symmetric key. Symmetric encryption encrypts the data.

Remember this

Knowing which key encrypts and which key decrypts will help you answer many questions on the exam. For example, just by knowing that a private key is encrypting, you know that it is being used for a digital signature.

Protecting Email

Cryptography provides two primary security methods you can use with email: digital signatures and encryption. These are separate processes, but you can digitally sign and encrypt the same email.

Signing Email with Digital Signatures

Digital signatures are similar in concept to handwritten signatures on printed documents that identify individuals, but they provide more security benefits. A digital signature is an encrypted hash of a message, encrypted with the sender's private key. If the recipient of a digitally signed email can decrypt the hash, it provides the following three security benefits:

- **Authentication.** This identifies the sender of the email. Email recipients have assurances the email actually came from who it appears to be coming from. For example, if an executive digitally signs an email, recipients know it came from the executive and not from an attacker impersonating the executive.
- **Non-repudiation.** The sender cannot later deny sending the message. This is sometimes required with online transactions. For example, imagine if Homer sends an order to sell stocks using a digitally signed email. If the stocks increase after his sale completes, he can't deny the transaction.
- **Integrity.** This provides assurances that the message has not been modified or corrupted. Recipients know that the message they received is the same as the sent message.

Digital signatures are much easier to grasp if you understand some other cryptography concepts discussed in this chapter. As a short review, these concepts are:

- **Hashing.** Digital signatures start by creating a hash of the message. A hash is simply a number created by performing an algorithm on the message.
- **Certificates.** Digital signatures need certificates, and certificates include the sender's public key.
- **Public/private keys.** In a digital signature, the sender uses the sender's private key to encrypt the hash of the message. The recipient uses the sender's public key to decrypt the hash of the message.

Figure 10.5 shows an overview of this process. In the figure, Lisa is sending a message to Bart with a digital signature. Note that the message "I passed" is not secret. If it was, Lisa would encrypt it, which is a completely separate process. The focus in this explanation is only the digital signature.

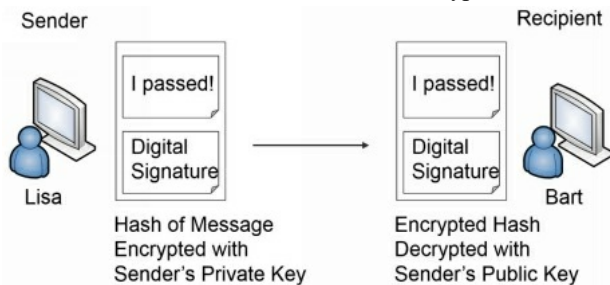


Figure 10.5: Digital signature process

Lisa creates her message in an email program, such as Microsoft Outlook. Once Microsoft Outlook is configured, all she has to do is click a button to digitally sign the message. Here is what happens when she clicks the button:

1. The application hashes the message.
2. The application retrieves Lisa's private key and encrypts the hash using this private key.
3. The application sends both the encrypted hash and the unencrypted message to Bart.

When Bart's system receives the message, it verifies the digital signature using the following steps:

1. Bart's system retrieves Lisa's public key, which is in Lisa's public certificate. In some situations, Lisa may have sent Bart a copy of her certificate with her public key. In domain environments, Bart's system can automatically retrieve Lisa's certificate from a network location.
2. The email application on Bart's system decrypts the encrypted hash with Lisa's public key.
3. The application calculates the hash on the received message.
4. The application compares the decrypted hash with the calculated hash.

If the calculated hash of the received message is the same as the encrypted hash of the digital signature, it validates several important checks:

- **Authentication.** Lisa sent the message. The public key can only decrypt something encrypted with the private key, and only Lisa has the private key. If the decryption succeeded, Lisa's private key must have encrypted the hash. On the other hand, if another key was used to encrypt the hash, Lisa's public key could not decrypt it. In this case, Bart will see an error indicating a problem with the digital signature.
- **Non-repudiation.** Lisa cannot later deny sending the message. Only Lisa has her private key and if the public key decrypted the hash, the hash must have been encrypted with her private key. Non-repudiation is valuable in online transactions.

- **Integrity.** Because the hash of the sent message matches the hash of the received message, the message has maintained integrity. It hasn't been modified.

Remember this

A digital signature is an encrypted hash of a message. The sender's private key encrypts the hash of the message to create the digital signature. The recipient decrypts the hash with the sender's public key. If successful, it provides authentication, non-repudiation, and integrity. Authentication identifies the sender. Integrity verifies the message has not been modified. Non-repudiation prevents senders from later denying they sent an email.

At this point, you might be thinking, if we do all of this, why not just encrypt the message, too? The answer is resources. It doesn't take much processing power to encrypt 256 bits in a SHA-256 hash. In contrast, it would take quite a bit of processing power to encrypt a lengthy email and its attachments. However, if you need to ensure confidentiality of the email, you can encrypt it.

Encrypting Email

There are times when you want to ensure that email messages are only readable by authorized users. You can encrypt email and just as any other time encryption is used, encrypting an email provides confidentiality.

Encrypting Email with Only Asymmetric Encryption

Imagine that Lisa wants to send an encrypted message to Bart. The following steps provide a simplified explanation of the process if only asymmetric encryption is used:

1. Lisa retrieves a copy of Bart's certificate that contains his public key.
2. Lisa encrypts the email with Bart's public key.
3. Lisa sends the encrypted email to Bart.
4. Bart decrypts the email with his private key.

This works because Bart is the only person who has access to his private key. If attackers intercepted the email, they couldn't decrypt it without Bart's private key. With this in mind, it's important to remember that when you're encrypting email contents, the recipient's public key encrypts and the recipient's private key decrypts. The sender's keys are not involved in this process. In contrast, a digital signature only uses the sender's keys but not the recipient's keys.

In most cases, the public key doesn't actually encrypt the message, but instead encrypts a symmetric key used to encrypt the email. The recipient then uses the private key to decrypt the symmetric key, and then uses the symmetric key to decrypt the email.

Remember this

The recipient's public key encrypts when encrypting an email message and the recipient uses the recipient's private key to decrypt an encrypted email message.

Encrypting Email with Asymmetric and Symmetric Encryption

The previous description provides a simplistic explanation of email encryption used by some email applications. However, most email applications combine both asymmetric and symmetric encryption. You may remember from earlier in this chapter that asymmetric encryption is slow and inefficient, but symmetric encryption is very quick.

Instead of using only symmetric encryption, most email applications use asymmetric encryption to privately share a session key. They then use symmetric encryption to encrypt the data. For example, imagine that Lisa is sending Bart an encrypted message. Figure 10.6 shows the process of encrypting the message and the symmetric key. Figure 10.7 shows the process of sending the encrypted message and encrypted session key, and identifies how the recipient can decrypt the data:

1. Lisa identifies a symmetric key to encrypt her email. For this example, assume it's a simplistic symmetric key of 53, though a symmetric algorithm like AES would use 128-bit or larger keys.
2. Lisa encrypts the email contents with the symmetric key of 53.
3. Lisa retrieves a copy of Bart's certificate that contains his public key.
4. She uses Bart's public key to encrypt the symmetric key of 53.
5. Lisa sends the encrypted email and the encrypted symmetric key to Bart.
6. Bart decrypts the symmetric key with his private key.
7. He then decrypts the email with the decrypted symmetric key.




Figure 10.6: Encrypting email

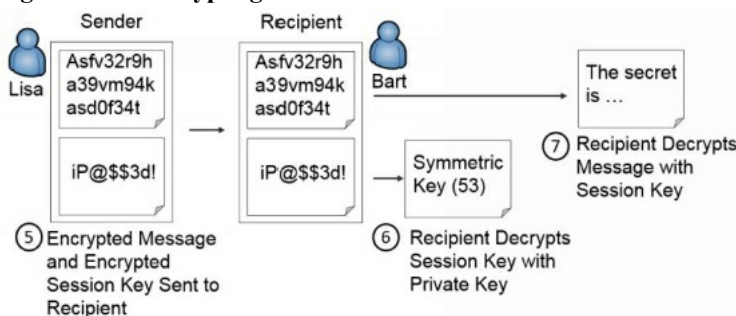


Figure 10.7: Decrypting email

Unauthorized users who intercept the email sent by Lisa won't be able to read it because it's encrypted with the symmetric key. Additionally, they can't read the symmetric key because it's encrypted with Bart's public key, and only Bart's private key can decrypt it.

S/MIME

Secure/Multipurpose Internet Mail Extensions (S/MIME) is one of the most popular standards used to digitally sign and encrypt email. Most email applications that support encryption and digital signatures use S/MIME standards.

S/MIME uses RSA for asymmetric encryption and AES for symmetric encryption. It can encrypt email at rest (stored on a drive) and in transit (data sent over the network). Because S/MIME uses RSA for asymmetric encryption, it requires a PKI to distribute and manage certificates.

PGP/GPG

Pretty Good Privacy (PGP) is a method used to secure email communication. It can encrypt, decrypt, and digitally sign email. Phillip Zimmerman designed PGP in 1991, and it has gone through many changes and improvements over the years and has been bought and sold by many different companies. Symantec Corporation purchased it in June 2010.

OpenPGP is a PGP-based standard created to avoid any conflict with existing licensing. In other words, users have no obligation to pay licensing fees to use it. Some versions of PGP follow S/MIME standards. Other versions follow OpenPGP standards. GNU Privacy Guard (GPG) is free software that is based on the OpenPGP standard.

Each of the PGP versions uses the RSA algorithm and public and private keys for encryption and decryption. Just like S/MIME, PGP uses both asymmetric and symmetric encryption.

Transport Encryption

Transport encryption methods encrypt data in transit to ensure transmitted data remains confidential. This includes data transmitted over the Internet and on internal networks. Earlier chapters cover several transport encryption protocols. As a reminder, they are:

- **SSH.** Chapter 3, “Understanding Basic Network Security,” presented Secure Shell (SSH). It’s used to encrypt a wide assortment of traffic, such as Secure File Transport Protocol (SFTP), Secure Copy (SCP), and Telnet. SSH uses port 22.
- **HTTPS.** Chapter 3 presented HTTP Secure. HTTPS uses either SSL or TLS to encrypt web traffic over port 443.

The primary transport encryption methods discussed in this section are IPsec, SSL, and TLS.

IPsec

Chapters 3 and 4 introduced IPsec. As a reminder, it can encrypt data in Tunnel mode with virtual private networks (VPNs) such as with L2TP/IPsec. It can also encrypt data in Transport mode between two systems.

IPsec uses an Authentication Header (AH) to provide authentication and integrity. Request for Comments 4835 (RFC 4835) mandates the use of HMAC for AH. Routers and firewalls identify AH traffic with protocol ID 51.

It can also use Encapsulating Security Payload (ESP) to provide confidentiality, in addition to the authentication and integrity provided by AH. RFC 4835 mandates the use of AES or 3DES as the symmetric encryption algorithm. Routers and firewalls identify ESP traffic with protocol ID 50.

When IPsec uses ESP, it encapsulates the entire IP packet and adds an additional IP header. The original IP header includes information such as destination ports, which can give clues to what type of data is in the payload. However, by encapsulating the original IP header and creating a new one, attackers can only see that the packet is IPsec with ESP.

Remember this

Transport encryption methods such as SSH, IPsec, HTTPS, SSL, and TLS protect the confidentiality of data transmitted over a network. IPsec must use HMAC for authentication and integrity. It can use either AES or 3DES for encryption with ESP. When IPsec uses ESP, it encrypts the entire packet, including the original IP header, and creates an additional IP header.

SSL

Secure Sockets Layer (SSL) is an encryption protocol used to encrypt Internet traffic. For example, HTTPS uses SSL in secure web browser sessions. It can also encrypt other transmissions. For example, File Transport Protocol Secure (FTPS) uses SSL to encrypt transmissions.

SSL provides certificate-based authentication and encrypts data with a combination of both symmetric and asymmetric encryption during a session. It uses asymmetric encryption to privately share a session key and symmetric encryption to encrypt data displayed on the web page and transmitted during the session.

Netscape created SSL for its web browser and updated it to version SSL 3.0. This was before organizations such as the Internet Engineering Task Force (IETF) created and maintained standards. Netscape’s success waned and there wasn’t a standardization process to update SSL, even though all web browsers were using it. The IETF created TLS to standardize improvements with SSL.

TLS

Transport Layer Security (TLS) is a replacement for SSL and is widely used in many different applications. The IETF has updated and published several TLS documents specifying the standard. TLS 1.0 is based on SSL 3.0 and is referred to as SSL 3.1. Similarly, each update to TLS indicates it is an update to SSL. For example, TLS 1.1 is called SSL 3.2 and TLS 1.2 is called SSL 3.3.

Just like SSL, TLS provides certificate-based authentication and uses both asymmetric and symmetric encryption. It uses asymmetric encryption to privately share a symmetric key and uses symmetric encryption to encrypt data in the web session. The “Encrypting HTTPS Traffic with SSL or TLS” section later in this chapter shows this process.

Many other applications use TLS. For example, Chapter 4 introduced the Extensible Authentication Protocol (EAP) in the context of increasing wireless security by adding authentication with an 802.1x server. Protected EAP (PEAP) and EAP-Tunneled TLS (EAP-TTLS) require 802.1x servers to have a certificate. EAP-TLS is the most secure method because it requires both servers and clients to use certificates. TLS can encrypt other traffic such as FTP and SMTP just as SSL does.

It’s important to remember that TLS and SSL require certificates. Certificate authorities (CAs) issue and manage certificates, so a CA is required to support TLS and SSL. These CAs can be internal or external third-party CAs.

Remember this

TLS is the replacement for SSL. Both TLS and SSL require certificates issued by certificate authorities (CAs). For example, PEAP-TLS uses TLS to encrypt the authentication process and requires a certificate

issued by a CA.

Cipher Suites

Cipher suites are a combination of cryptographic algorithms that provide several layers of security for TLS and SSL. When two systems connect, they identify a cipher suite that is acceptable to both systems and then use the protocols within that suite. The protocols within the suite provide three primary cryptographic solutions. They are:

- **Encryption.** Encryption provides confidentiality of data. TLS and SSL use asymmetric cryptography to privately exchange a symmetric key and then encrypt the data with a symmetric algorithm. TLS and SSL support several types of symmetric encryption, including RC4, 3DES, and AES.
- **Authentication.** TLS and SSL use certificates for authentication. Clients can verify the authenticity of the certificate by querying the CA that issued the certificate.
- **Integrity.** TLS and SSL use a message authentication code (MAC) for integrity. For example, they can use HMAC-MD5 or HMAC-SHA1.

There are over 200 named cipher suites, and systems identify them with a cipher identifier as a string of hexadecimal characters and a coded name.

Here are two examples:

- **0x00C031.** TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256
- **0x00003C.** TLS_RSA_WITH_AES_128_CBC_SHA256

If you're familiar with the acronyms, you can get a good idea of what each cipher suite is using. Here are some notes for clarification:

- **Protocol.** Both are using TLS. If they were SSL, the first three letters would be SSL or SSL2.
- **Key exchange method.** The first one is using ECDH and the second one is using RSA.
- **Authentication.** Both are using RSA, though, they shortened the code in the second one. Instead of listing RSA twice, it is only listed once.
- **Encryption.** Both are using 128-bit AES, though in different modes. Galois/Counter Mode (GCM) and Cipher-Block Chaining (CBC) are the two modes identified here. You don't need to know the modes for the CompTIA Security+ exam. However, if you want to dig into them, check out SP 800-38 A through SP 800-38 F, which you can download from <http://csrc.nist.gov/publications/PubsSPs.html>.
- **Integrity.** Both are using the SHA-256 hashing algorithm.

Some cipher suites are very old and include encryption algorithms such as DES. Clearly, they shouldn't be used and are disabled by default in most systems today. When necessary, administrators configure systems and applications to disable older specific cipher suites.

When two systems connect, they negotiate to identify which cipher suite they use. In essence, each system passes a prioritized list of cipher suites it is willing to use. They then use the cipher suite that is the highest on each list.

Strong Versus Weak Ciphers

If security was the only concern, every connection would use the strongest cipher suite all the time. However, that isn't feasible. Due to hardware and software restrictions, some systems do not support some of the stronger ciphers. Additionally, all data isn't the same. You would use the strongest cipher to encrypt Top Secret data, but this isn't needed to encrypt data with a lower classification level.

With this in mind, administrators configure systems to use specific cipher suites based on their needs. Imagine you need to configure transport encryption to protect data transferred to and from a server hosting financial data. What would you choose? There are many choices.

First, you would want to use TLS instead of SSL because it is an update to SSL and TLS is stronger. In contrast, if you configuring a web site to support HTTPS, you would probably include support for both TLS and SSL with a preference for TLS. This ensures all clients can still connect, even if their systems do not support TLS.

Additionally, you would use a strong encryption algorithm such as AES-256. You would also use a strong hashing algorithm such as HMAC-SHA1.

Encrypting HTTPS Traffic with SSL or TLS

HTTP Secure (HTTPS) is commonly used on the Internet to secure web traffic. HTTPS can use either SSL or TLS to encrypt the traffic, and both use asymmetric and symmetric encryption. If you're able to grasp the basics of how HTTPS combines both asymmetric and symmetric encryption, you'll have what you need to know for most protocols that use both encryptions.

Because asymmetric encryption isn't efficient to encrypt large amounts of data, symmetric encryption is used to encrypt the session data. However, both the client and the server must know what this symmetric key is before they can use it. They can't whisper it to each other over the Internet. That's like an actor on TV using a loud whisper, or stage whisper, to share a secret. Millions of TV viewers can also hear the secret.

Instead, HTTPS uses asymmetric encryption to securely transmit a symmetric key. It then uses the symmetric key with symmetric encryption to encrypt all the data in the HTTPS session.

Figure 10.8 and the following steps show the overall process of establishing and using an HTTPS session. As you read these steps, try to keep these two important concepts in mind:

- SSL and TLS use *asymmetric* encryption to securely share the symmetric key.
- SSL and TLS use *symmetric* encryption to encrypt the session data.

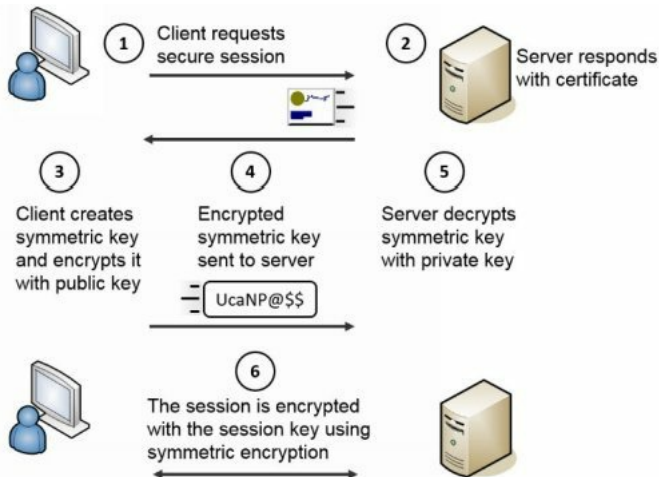


Figure 10.8: Simplified handshake process used with HTTPS

1. The client begins the process by requesting an HTTPS session. This could be by entering an HTTPS address in the URL or by clicking on an HTTPS link.
2. The server responds by sending the server's certificate. The certificate includes the server's public key. The matching private key is on the server and only accessible by the server.
3. The client creates a symmetric key and encrypts it with the server's public key.
As an example, imagine that the symmetric key is 53 (though in reality it would be much more complex). The client encrypts the session key of 53 using the web server's public key creating ciphertext of UcaNP@\$\$.
This symmetric key will be used to encrypt data in the HTTPS session, so it is sometimes called a session key.
4. The client sends the encrypted session key (UcaNP@\$) to the web server. Only the server's private key can decrypt this. If attackers intercept the encrypted key, they won't be able to decrypt it because they don't have access to the server's private key.
5. The server receives the encrypted session key and decrypts it with the server's private key. At this point, both the client and the server know the session key.
6. All of the session data is encrypted with this symmetric key using symmetric encryption.

The amazing thing to me is that this happens so quickly. If a web server takes as long as five seconds, many of us wonder why it's taking so long. However, a lot is happening to establish this session.

Key Stretching

Key stretching is a technique used to increase the strength of stored passwords. Chapter 7 discusses many types of password attacks, including brute force attacks and rainbow table attacks. Using complex passwords goes a long way toward preventing these types of attacks but doesn't prevent them all. A better method is to salt the passwords with additional random bits to make them even more complex and key stretching techniques use salts.

Two common key stretching techniques are bcrypt and Password-Based Key Derivation Function 2 (PBKDF2):

- **Bcrypt.** Based on the Blowfish block cipher, bcrypt is used on many Unix and Linux distributions to protect the passwords stored in the shadow password file. Bcrypt salts the password by adding additional bits before encrypting it with Blowfish.
- **PBKDF2.** Many algorithms such as Wi-Fi Protected Access II (WPA2), Apple's iOS mobile operating system, and Cisco operating systems use PBKDF2 to increase the security of passwords. PBKDF2 adds a salt of at least 64 bits.

Remember this

Bcrypt and PBKDF2 are key stretching techniques that help prevent brute force and rainbow table attacks. Both salt the password with additional bits.

In-Band Versus Out-of-Band Key Exchange

In-band key exchange indicates that the two parties share an encryption key in the same communication channel as the encrypted data. This can be risky because anyone who captures the exchange will have the key and can decrypt the data. Out-of-band key exchange indicates that the two parties share the symmetric key in one communication channel and then exchange the encrypted data in a separate communication channel.

The previous sections covered several asymmetric encryption topics. As a reminder, the primary purpose of asymmetric encryption is to privately share a session key used for symmetric encryption. Another way of saying this is that asymmetric encryption methods share the session key using an out-of-band key exchange method.

Exploring PKI Components

A Public Key Infrastructure (PKI) is a group of technologies used to request, create, manage, store, distribute, and revoke digital certificates.

Asymmetric encryption depends on the use of certificates for a variety of purposes, such as protecting email and protecting Internet traffic with SSL and TLS. For example, HTTPS sessions protect Internet credit card transactions, and these transactions depend on a PKI.

A primary benefit of a PKI is that it allows two people or entities to communicate securely without knowing each other previously. In other words, it allows them to communicate securely through an insecure public medium such as the Internet.

For example, you can establish a secure session with Amazon.com even if you’ve never done so before. Amazon purchased a certificate from VeriSign. As shown in the “Encrypting HTTPS Traffic with SSL or TLS” section previously, the certificate provides the ability to establish a secure session.

A key element in a PKI is a Certificate Authority.

Certificate Authority

A Certificate Authority (CA, pronounced “cah”) issues, manages, validates, and revokes certificates. In some contexts, you might see a CA referred to as a certification authority, but they are the same thing. CAs can be very large, such as VeriSign, which is a public CA. A CA can also be very small, such as a single service running on a server in a domain.

Public CAs make money by selling certificates. For this to work, the public CA must be trusted. Certificates issued by the CA are trusted as long as the CA is trusted.

This is similar to how a driver’s license is trusted. The Department of Motor Vehicles (DMV) issues driver’s licenses after validating a person’s identity. If you want to cash a check, you may present your driver’s license to prove your identity. Businesses trust the DMV, so they trust the driver’s license. On the other hand, if you purchased an ID from Gibson’s Instant IDs, businesses may not trust it.

Although we may trust the DMV, why would a computer trust a CA? The answer is based on the certificate trust path.

Certificate Trust Paths and Trust Models

CAs are trusted by placing a copy of their root certificate into a trusted root CA store. The root certificate is the first certificate created by the CA that identifies it, and the store is just a collection of these root certificates. If the CA's root certificate is placed in this store, all certificates issued by this CA are trusted.

Figure 10.9 shows the Trusted Root Certification Authority store from Internet Explorer. You can see that there are many certificates from many different CAs. VeriSign is a popular CA, so I scrolled down to show root certificates from VeriSign.

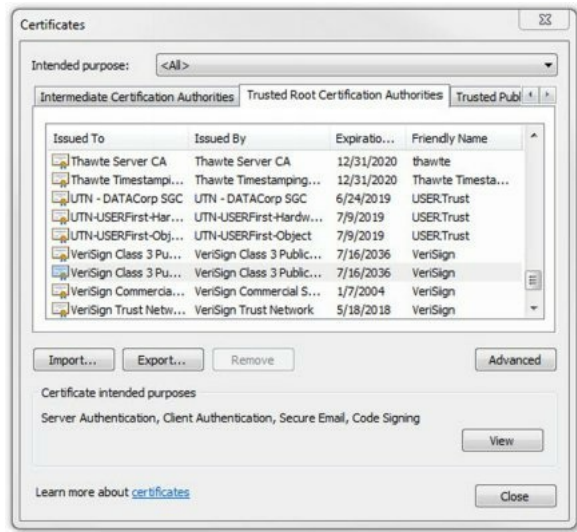


Figure 10.9: Trusted Root Certification Authorities

Public CAs such as VeriSign negotiate with web browser developers to have their certificates included with the web browser. This way, any certificates that they sell to businesses are automatically trusted.

The most common trust model is the hierarchical trust model, also known as a centralized trust model. In this model, the public CA creates the first CA, known as the root CA. If the organization is large, it can create child CAs. A large trust chain works like this:

- The root CA issues certificates to intermediate CAs.
- Intermediate CAs issue certificates to child CAs.
- Child CAs issue certificates to devices or end users.

In a small organization, the root CA can simply issue certificates to the devices and end users. It's not necessary to have intermediate and child CAs.

Another type of trust model is a web of trust or decentralized trust model, sometimes used with PGP and GPG. A web of trust uses self-signed certificates, and a third party vouches for these certificates. For example, if five of your friends trust a certificate, you can trust the certificate. If the third party is a reliable source, the web of trust provides a secure alternative. However, if the third party does not adequately verify certificates, it can result in the use of certificates that shouldn't be trusted.

Self-Signed Certificates

It is possible to create a CA and use self-signed certificates. For example, an administrator can use Active Directory Certificate Services (AD CS) on Windows Server 2008 to create a CA and issue certificates to company-owned web servers. AD CS is built in to Windows Server 2008, so this is certainly less expensive than purchasing a certificate from a public CA.

However, certificates issued by this CA will not be trusted by default. If a user connects to this web server and establishes an HTTPS session, the web browser will show an error. Depending on the web browser, it may indicate that the issuer of a certificate (the CA) is not recognized or indicate that the site's certificate is not trusted. The error is often accompanied with warning icons and other notes encouraging the users not to continue.

This is not acceptable for an e-commerce web site. Imagine if Lisa has her credit card in hand ready to buy a product, and then sees errors indicating trust problems. She very likely won't continue.

If only employees use this web site, they could ignore the errors and click through to establish the connection. However, this may breed complacency. When they go to a public site and see this same type of error, they may ignore the error and continue.

Instead, an administrator can copy the CA's root certificate to the user's computer. Web browsers will then trust the certificate from the company web site, eliminating the errors.

Wildcard Certificates

In some situations, organizations choose to use wildcard certificates to reduce the management burden associated with certificates. As an example, imagine a web site with the domain name of *GetCertifiedGetAhead.com*. If a typical certificate is associated with this web site, it will only be valid for web pages in that domain. However, an organization can get a wildcard certificate for *. *GetCertifiedGetAhead.com*. This wildcard certificate is valid for the following domain names, too:

- *train.GetCertifiedGetAhead.com*
- *blogs.GetCertifiedGetAhead.com*

The * in the wildcard certificate will work for any single name added onto the domain name. In other words, administrators only need to purchase and install a single certificate for all three domain names. However, it only works for one level of domain names. It wouldn't be valid for something like this: *north.train.GetCertifiedGetAhead.com*.

Remember this

A PKI requires a trust model between CAs. Most trust models are hierarchical and centralized with a central root CA. Wildcard certificates reduce the management burden associated with certificates.

Registration

Users and systems request certificates from a CA using a registration process. In some cases, a user enters information manually into a web site form. In other cases, a user sends a specifically formatted file to the CA. Within a domain, the system handles much of the process automatically.

As an example, imagine I wanted to purchase a certificate for *GetCertifiedGetAhead.com* for secure HTTPS sessions. I would first create a public and private key pair. Many programs are available to automate this process. For example, OpenSSL is a command-line program that is included in many Linux distributions. It creates key pairs in one command and allows you to export the public key to a file in a second command. Technically, OpenSSL and similar applications create the private key first. However, these applications appear to create both keys at the same time.

I would then put together a certificate signing request (CSR) for the certificate, including the purpose of the certificate and information about the web site, the public key, and me. Most CAs require CSRs to be formatted using the public-key cryptography standards (PKCS) #10 specification.

I then send the CSR to the CA. The CA validates my identity and creates a certificate with the public key. The validation process is different based on the usage of the certificate. In some cases, it includes extensive checking, and in other cases, verification comes from the credit card I use to pay.

I can then register this certificate with my web site along with the private key. Any time someone initiates a secure HTTPS connection, the web site sends the certificate with the public key and the TLS/SSL session creates the session.

In large organizations, a registration authority (RA) can assist the CA by collecting registration information. The RA never issues certificates. Instead, it only assists in the registration process.

Remember this

You typically request certificates using a certificate signing request (CSR). The first step is to create the RSA-based private key, which is used to create the public key. You then include the public key in the CSR and the CA will embed the public key in the certificate.

Revoking Certificates

Normally, certificates expire based on the Valid From and Valid To dates. However, there are some instances when a CA will revoke a certificate before it expires.

For example, if a private key is publicly available, the key pair is compromised. It no longer provides adequate security because the private key is no longer private. Similarly, if the CA itself is compromised through a security breach, certificates issued by the CA may be compromised, so the CA can revoke certificates.

In general, any time a CA does not want anyone to use a certificate, the CA revokes it. Although the most common reasons are due to compromise of a key or the CA, there are others. A CA can use any of the following reasons when revoking a certificate:

- Key compromise
- CA compromise
- Change of affiliation
- Superseded
- Cease of operation
- Certificate hold

CAs use certificate revocation lists (CRL, pronounced “crill”) to revoke a certificate. The CRL is a version 2 certificate that includes a list of revoked certificates by serial number. For example, Figure 10.10 shows a copy of a CRL.

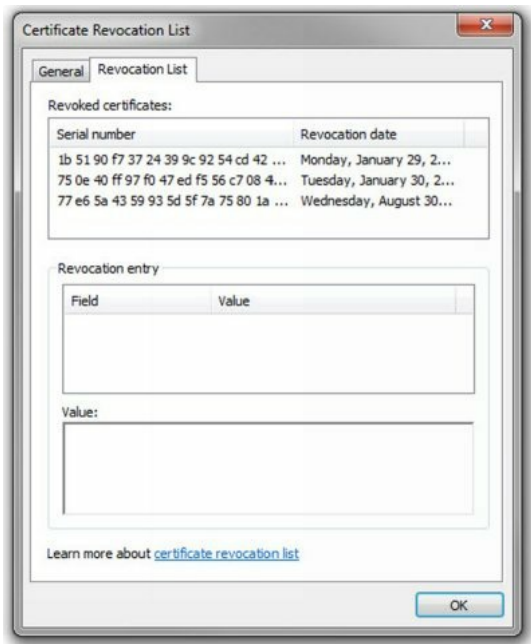


Figure 10.10: Certificate revocation list

Validating Certificates

Before clients use a certificate, they first verify it is valid. The first check is to ensure that it isn't expired. If the certificate is expired, the computer system typically gives the user an error indicating the certificate is not valid.

Clients also validate certificates through the CA. First, they verify that the certificate was issued by a trusted CA. Next, they query the CA to verify the CA hasn't revoked the certificate. A common method of validating a certificate is by requesting a copy of the CRL, as shown in Figure 10.11. The following steps outline the process:

1. The client initiates a session requiring a certificate, such as an HTTPS session.
2. The server responds with a copy of the certificate that includes the public key.
3. The client queries the CA for a copy of the CRL.
4. The CA responds with a copy of the CRL.




Figure 10.11: Validating a certificate

The client then checks the serial number of the certificate against the list of serial numbers in the CRL. If the certificate is revoked for any reason, the application gives an error message to the user.

Notice that the CA and the web server are both protected within demilitarized zones (DMZs). As mentioned in Chapter 3, a DMZ is a buffered zone between an internal network and the Internet. The DMZ provides a layer of protection for Internet-facing servers, but servers in the DMZ are available on the Internet.

Another method of validating a certificate is with the Online Certificate Status Protocol (OCSP). OCSP allows the client to query the CA with the serial number of the certificate. The CA then responds with an answer of “good,” “revoked,” or “unknown.” A response of “unknown” could indicate the certificate is a forgery. OCSP can be more efficient, especially if the CRL includes a large number of revoked certificates.

Remember this

CAs revoke certificates for several reasons such as when the private key is compromised or the CA is compromised. An internal CA can revoke a certificate when an employee leaves. The certificate revocation list (CRL) includes a list of revoked certificates and is publicly available. An alternative to using a CRL is the Online Certificate Status Protocol (OCSP), which returns answers such as good, revoked, or unknown.

Outdated Certificates

In some situations, users can find that they can no longer access data because they don't have access to a previously used certificate. The classic example occurs when users are issued new smart cards. As an example, many contractors have contracts of a specific length, such as one year. When the year expires, the contract is typically renewed, but they need to get a new smart card. They'll have problems digitally signing email and opening encrypted email.

The problem is that their previous certificate is published in the organization, typically in a global address list accessible to everyone else in the organization. What these users need to do is publish the new certificate embedded in their new smart card and this will resolve the problem.

Remember this

When an organization issues a user a new certificate, such as with a new smart card, the user needs to publish it within the organization, such as via a global address list. Otherwise, other people in the organization use the original certificate.

Key Escrow

Key escrow is the process of placing a copy of a private key in a safe environment. This is useful for recovery. If the original is lost, the organization retrieves the copy of the key to access the data. Key escrow isn't required, but if an organization determines that data loss is unacceptable, it will implement a key escrow process.

In some cases, an organization provides a copy of the key to a third party. Another method is to designate employees within the organization who will be responsible for key escrow. These employees maintain and protect copies of the key, and if the original key is lost, they check out a copy of the key to an administrator or user.

Recovery Agent

A key recovery agent is a designated individual who can recover or restore cryptographic keys. In the context of a PKI, a recovery agent can recover private keys to access encrypted data. The recovery agent may be a security professional, administrator, or anyone designated by the company.

In some cases, the recovery agent can recover encrypted data using a different key. For example, Microsoft's BitLocker supports encryption of entire drives. It's possible to add a data recovery agent field when creating a BitLocker encrypted drive. In this case, BitLocker uses two keys. The user has one key and uses it to unlock the drive during day-to-day use. The second key is only accessible by the recovery agent and used for recovery purposes if the original key is lost or becomes inaccessible.

Remember this

Recovery agents can recover user messages and data when users lose access to their private keys. In some cases, recovery agents can recover the private key from a key escrow.

Chapter 10 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Providing Integrity with Hashing

- Hashing verifies the integrity of data, such as downloaded files and email messages.
- A hash (sometimes listed as a checksum) is a fixed-size string of numbers or hexadecimal characters.
- Hashing algorithms are one-way functions used to create a hash. You cannot reverse the process to re-create the original data.
- Passwords are often stored as hashes instead of the actual password. Salting the password thwarts many password attacks.
- Common hashing algorithms are Message Digest 5 (MD5), Secure Hash Algorithm (SHA), and Hash-based Message Authentication Code (HMAC). HMAC provides both integrity and authenticity of a message.
- Transport encryption protocols such as Internet Protocol security (IPsec) and Transport Layer Security (TLS) use HMAC-MD5 and HMAC-SHA1.

Providing Confidentiality with Encryption

- Confidentiality ensures that data is only viewable by authorized users. Encryption provides confidentiality of data, including data at rest (any type of data stored on disk) and data in transit (any type of transmitted data).
- Symmetric encryption uses the same key to encrypt and decrypt data. As an example, Remote Authentication Dial-In User Service (RADIUS) uses a shared key for symmetric encryption.
- Block ciphers encrypt data in fixed-size blocks. Advanced Encryption Standard (AES) and Twofish encrypt data in 128-bit blocks.
- Stream ciphers encrypt data one bit or one byte at a time. They are more efficient than block ciphers when encrypting data of an unknown size, or sent in a continuous stream. RC4 is a commonly used stream cipher.
- Data Encryption Standard (DES), Triple DES (3DES), and Blowfish are block ciphers that encrypt data in 64-bit blocks.
- AES is a popular symmetric block encryption algorithm, and it uses 128, 192, or 256 bits for the key.
- DES is an older, symmetric block encryption algorithm. 3DES was created as an improvement over DES and is used when hardware doesn't support AES.
- One-time pads provide the strongest encryption when compared with other encryption methods.
- Asymmetric encryption uses public and private keys as matched pairs.
 - If the public key encrypted information, only the matching private key can decrypt it.
 - If the private key encrypted information, only the matching public key can decrypt it.
 - Private keys are always kept private and never shared.
 - Public keys are freely shared by embedding them in a certificate.
- RSA is a popular asymmetric algorithm. Many cryptographic protocols use RSA to secure data such as email and data transmitted over the Internet. RSA uses prime numbers to generate public and private keys.
- Elliptic curve cryptography (ECC) is an encryption technology commonly used with small wireless devices.
- Diffie-Hellman provides a method to privately share a symmetric key between two parties. Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) is a version of Diffie-Hellman that uses ECC to re-create keys for each session.
- Steganography is the practice of hiding data within a file. You can hide messages in the white space of a file without modifying its size. A more sophisticated method is by modifying bits within a file. Capturing and comparing hashes of files can discover steganography attempts.
- Transport encryption methods protect the confidentiality of data sent over the network. IPsec, TLS, and SSL are three examples.
- IPsec uses HMAC for authentication and integrity and AES or 3DES for encryption.
- TLS is the replacement for SSL. Both require certificates issued from a CA.

Using Cryptographic Protocols

- When using digital signatures with email:
 - The sender's private key encrypts (or signs).
 - The sender's public key decrypts.
- A digital signature provides authentication (verified identification) of the sender, non-repudiation, and integrity of the message.
 - Senders create a digital signature by hashing a message and encrypting the hash with the sender's private key.
 - Recipients decrypt the digital signature with the sender's matching public key.
- When encrypting email:
 - The recipient's public key encrypts.

- The recipient's private key decrypts.
- Many email applications use the public key to encrypt a symmetric key, and then use the symmetric key to encrypt the email contents.
- When encrypting web site traffic with SSL or TLS:
 - The web site's public key encrypts a symmetric key.
 - The web site's private key decrypts the symmetric key.
 - The symmetric key encrypts data in the session.
- S/MIME and PGP secure email with encryption and digital signatures. They both use RSA, certificates, and depend on a PKI. They can encrypt email at rest (stored on a drive) and in transit (sent over the network).
- Two commonly used key stretching techniques are bcrypt and Password-Based Key Derivation Function 2 (PBKDF2). They protect passwords against brute force and rainbow table attacks.

Exploring PKI Components

- A Public Key Infrastructure (PKI) is a group of technologies used to request, create, manage, store, distribute, and revoke digital certificates. A PKI allows two entities to privately share symmetric keys without any prior communication.
- Most public CAs use a hierarchical centralized CA trust model, with a root CA and intermediate CAs.
- A CA issues, manages, validates, and revokes certificates. Wildcard certificates use a * for child domains to reduce the administrative burden of managing certificates.
- Root certificates of trusted CAs are stored on computers. If a CA's root certificate is not in the trusted store, web users will see errors indicating the certificate is not trusted or the CA is not recognized.
- You request a certificate with a certificate signing request (CSR). You first create a private/public key pair and include the public key in the CSR.
- CAs revoke certificates when an employee leaves, the private key is compromised, or the CA is compromised. A CRL identifies revoked certificates as a list of serial numbers.
- The CA publishes the CRL, making it available to anyone. Web browsers can check certificates they receive from a web server against a copy of the CRL to determine if a received certificate is revoked.
- User systems return errors when a system tries to use an expired certificate.
- When users are issued new certificates, such as in a new smart card, they need to publish the new certificate. This is typically done by publishing it to a global address list.
- A key escrow stores a copy of private keys used within a PKI. If the original private key is lost or inaccessible, the copy is retrieved from escrow, preventing data loss.
- Recovery agents can recover data secured with a private key, or recover a private key, depending on how the recovery agent is configured.

Chapter 10 Practice Questions

1. Of the following choices, what can you use to verify data integrity?
 - A. AES
 - B. DES
 - C. RC4
 - D. SHA
2. A security technician runs an automated script every night designed to detect changes in files. Of the following choices, what are the most likely protocols used in this script?
 - A. PGP and MD5
 - B. ECC and HMAC
 - C. AES and Twofish
 - D. MD5 and HMAC
3. Some encryption algorithms use stream ciphers and some use block ciphers. Which of the following are examples of block ciphers? (Choose THREE.)
 - A. AES
 - B. DES

- C. MD5
- D. SHA
- E. RC4
- F. Blowfish

4. Which of the following algorithms encrypts data in 64-bit blocks?

- A. AES
- B. DES
- C. Twofish
- D. RC4

5. An application developer needs to use an encryption protocol to encrypt credit card data within a database used by the application. Which of the following would be the FASTEST, while also providing strong confidentiality?

- A. AES-256
- B. DES
- C. Blowfish
- D. SHA-2

6. Your organization uses several different types of cryptographic techniques. Which of the following techniques uses a private key and a public key?

- A. AES
- B. RSA
- C. Blowfish
- D. MD5

7. Your network requires a secure method of sharing encryption keys over a public network. Which of the following is the BEST choice?

- A. Symmetric encryption
- B. Bcrypt
- C. Diffie-Hellman
- D. Steganography

8. Your organization plans to issue some employees mobile devices such as smartphones and tablets. These devices don't have a lot of processing power. Which of the following cryptographic methods has the LEAST overhead and will work with these mobile devices?

- A. ECC
- B. 3DES
- C. Bcrypt
- D. PBKDF2

9. A manager is suspected of leaking trade secrets to a competitor. A security investigator is examining his laptop and notices a large volume of vacation pictures on the hard drive. Data on this laptop automatically uploads to a private cloud owned by the company once a week. The investigator noticed that the hashes of most of the pictures on the hard drive are different from the hashes of the pictures in the cloud location. Which of the following is the MOST likely explanation for this scenario?

- A. The manager is leaking data using hashing methods.
- B. The manager is leaking data using digital signatures.
- C. The manager is leaking data using steganography methods.
- D. The manager is not leaking data.

10. A heavily used application accesses a financial database on a server within your network. Due to recent data breaches, management wants to ensure transport encryption protects this data. Which of the following algorithms is the BEST choice to meet this goal?

- A. SSL
- B. SHA
- C. TLS
- D. CRL

11. You are planning to encrypt data in transit. Which of the following protocols meets this need and encapsulates IP packets within an additional IP header?
- A. TLS
 - B. SSL
 - C. HMAC
 - D. IPsec
12. Homer wants to send a secure email to Marge so he decides to encrypt it. Homer wants to ensure that Marge can verify that he sent it. Which of the following does Marge need to verify the certificate that Homer used in this process is valid?
- A. The CA's private key
 - B. The CA's public key
 - C. Marge's public key
 - D. Marge's private key
13. Bart wants to send a secure email to Lisa so he decides to encrypt it. Bart wants to ensure that Lisa can verify that he sent it. Which of the following does Lisa need to meet this requirement?
- A. Bart's public key
 - B. Bart's private key
 - C. Lisa's public key
 - D. Lisa's private key
14. Users in your organization sign their emails with digital signatures. What provides integrity for these certificates?
- A. Hashing
 - B. Encryption
 - C. Non-repudiation
 - D. Private key
15. An application requires users to log on with passwords. The application developers want to store the passwords in such a way that it will thwart rainbow table attacks. Which of the following is the BEST solution?
- A. SHA
 - B. Blowfish
 - C. ECC
 - D. Bcrypt
16. Homer wants to use digital signatures for his emails and realizes he needs a certificate. Which of the following will issue Homer a certificate?
- A. CRL
 - B. CA
 - C. OCSP
 - D. Recovery agent
17. You need to submit a CSR to a CA. Which of the following would you do FIRST?
- A. Generate a new RSA-based session key.
 - B. Generate a new RSA-based private key.
 - C. Generate the CRL.
 - D. Implement OCSP.
18. Your organization is planning to implement an internal PKI. What is required to ensure users can validate certificates?
- A. An intermediate CA
 - B. CSR
 - C. Wildcard certificates
 - D. CRL

19. Your organization requires the use of a PKI and it wants to implement a protocol to validate trust with minimal traffic. Which of the following protocols validates trust by returning short responses, such as “good” or “revoked”?
- A. OCSP
 - B. CRL
 - C. CA
 - D. CSR
20. A user’s laptop developed a problem and can no longer boot. Help desk personnel tried to recover the data on the disk, but the disk is encrypted. Which of the following can be used to retrieve data from the hard drive?
- A. A trust relationship
 - B. Public key
 - C. Recovery agent
 - D. CRL

Chapter 10 Practice Question Answers

1. **D.** Secure Hash Algorithm (SHA) is one of many available hashing algorithms used to verify data integrity. None of the other options are hashing algorithms. Advanced Encryption Standard (AES), Data Encryption Standard (DES), and Rivest Cipher 4 (RC4) are symmetric encryption algorithms.
2. **D.** Hashing algorithms can detect changes in files (or verify the files have not lost integrity) and Message Digest 5 (MD5) and Hash-based Message Authentication Code (HMAC) are both hashing algorithms. Pretty Good Privacy (PGP) is a method used to secure email communication. Elliptic curve cryptography (ECC), Advanced Encryption Standard (AES), and TwoFish are all encryption algorithms.
3. **A, B, F.** Advanced Encryption Standard (AES), Data Encryption Standard (DES), and Blowfish are all block ciphers. Although it’s not listed, Triple DES (3DES) is also a block cipher. Message Digest 5 (MD5) and Secure Hash Algorithm (SHA) are hashing algorithms. Rivest Cipher 4 (RC4) is a stream cipher.
4. **B.** Data Encryption Standard (DES) encrypts data in 64-bit blocks similar to how 3DES and Blowfish encrypt data in 64-bit blocks. Advanced Encryption Standard (AES) and Twofish encrypt data in 128-bit blocks. Rivest Cipher 4 (RC4) is a stream cipher and it encrypts data one bit at a time.
5. **C.** Blowfish would be the fastest in this scenario. Blowfish provides strong encryption so would provide strong confidentiality. Advanced Encryption Standard-256 (AES-256) is a strong encryption protocol, but Blowfish is faster than AES in some situations such as when comparing it against AES-256. Data Encryption Standard (DES) is not secure and is not recommended today. Secure Hash Algorithm version 2 (SHA-2) is a hashing algorithm used for integrity.
6. **B.** Rivest, Shamir, Adleman (RSA) is an asymmetric algorithm and all asymmetric algorithms use public and private keys. Advanced Encryption Standard (AES) and Blowfish are strong block-based symmetric encryption algorithms. Message Digest 5 (MD5) is a hashing algorithm.
7. **C.** Diffie-Hellman allows entities to negotiate encryption keys securely over a public network. Once the entities negotiate the keys, they use symmetric encryption, but they can’t share keys using symmetric encryption without first using a secure method such as Diffie-Hellman. Bcrypt is a key stretching technique used by some Unix systems to make password cracking more difficult. Steganography hides data within data, but it isn’t the best method of sharing encryption keys over a public network.
8. **A.** Elliptic curve cryptography (ECC) has minimal overhead and is often used with mobile devices for encryption. Triple Data Encryption Standard (3DES) consumes a lot of processing time and isn’t as efficient as ECC. Password-Based Key Derivation Function 2 (PBKDF2) and bcrypt are key stretching techniques that salt passwords with additional bits to protect against brute force attempts.
9. **C.** The manager is most likely leaking data using steganography methods by embedding the data into the vacation pictures. If the file is the same, the hash of the file and the hash of a file copy should be the same. Because the hashes are different, it indicates the files are different and the most likely explanation is because some of the files have other data embedded within them. Hashing and digital signatures are not methods that would support leaking data. The scenario indicates the manager is suspected of leaking data, and the different hashes provide evidence to support this suspicion.
10. **C.** Transport Layer Security (TLS) is a transport encryption protocol that can protect the data while it is in transit. Secure Sockets Layer (SSL) is also a transport encryption protocol, but TLS is recommended instead. Secure Hash Algorithm (SHA) is a hashing algorithm, not an encryption protocol. Both SSL and TLS use certificates and revoked certificates are published in a certificate revocation list (CRL), but a CRL is not a transport encryption protocol.
11. **D.** Internet Protocol security (IPsec) can encrypt data in transit and encapsulates IP packets with an additional IP header. Transport Layer Security (TLS) and Secure Sockets Layer (SSL) are both transport encryption protocols that can protect the data while it is in transit. Although they both use certificates for security, they do not encapsulate IP packets within an additional IP header. Hash-based Message Authentication Code (HMAC) is often

used with IPsec, but HMAC does not encrypt data.

12. **B.** Marge would verify Homer's certificate is valid by querying the Certificate Authority (CA) that issued Homer's certificate and the CA's public certificate includes the CA's public key. Homer would use a digital signature to provide verification that he sent the message. Homer would encrypt the digital signature with his private key and Marge would decrypt the digital signature with Homer's public key. The CA's private key remains private. Marge's keys are not used for Homer's digital signature, but might be used for the encryption of the email.

13. **A.** Lisa would decrypt the digital signature with Bart's public key and verify the public key is valid by querying a Certificate Authority (CA). The digital signature provides verification that Bart sent the message, non-repudiation, and integrity for the message. Bart encrypts the digital signature with his private key, which can only be decrypted with his public key. Lisa's keys are not used for Bart's digital signature, but might be used for the encryption of the email. Although not part of this scenario, Bart would encrypt the email with Lisa's public key and Lisa would decrypt the email with Lisa's private key.

14. **A.** Hashing provides integrity for digital signatures and other data. A digital signature is a hash of the message encrypted with the sender's private key, but the encryption doesn't provide integrity. The digital signature provides non-repudiation, but non-repudiation does not provide integrity. The private key and public key are both needed, but the private key does not provide integrity.

15. **D.** Bcrypt is a key stretching technique designed to protect against brute force attempts and is the best choice of the given answers. Another alternative is Password-Based Key Derivation Function 2 (PBKDF2). Both salt the password with additional bits. Passwords stored using Secure Hash Algorithm (SHA) are easier to crack because they don't use salts. PBKDF2 is based on Blowfish, but Blowfish itself isn't commonly used to encrypt passwords. Elliptic curve cryptography (ECC) is efficient and sometimes used with mobile devices, but not to encrypt passwords.

16. **B.** A Certificate Authority (CA) issues and manages certificates. A certificate revocation list (CRL) is a list of revoked certificates. Online Certificate Status Protocol (OCSP) is an alternative to a CRL and validates certificates with short responses such as good, unknown, or revoked. A recovery agent can retrieve a private key if the original private key is no longer accessible.

17. **B.** You create the RSA-based private key first and then create the matching public key from it, which you include in the certificate signing request (CSR) that you send to the Certificate Authority (CA). The RSA algorithm technically creates the private key first, but most applications that create the key pair appear to create them at the same time. A session key is a symmetric key, but RSA is an asymmetric algorithm. The CA generates the certificate revocation list (CRL) to identify revoked certificates. Online Certificate Status Protocol (OCSP) is an alternative to using CRLs to validate certificates, but it is not required.

18. **D.** A certificate revocation list (CRL) includes a list of revoked certificates and it allows users to validate certificates. Any CA can issue a CRL, so an intermediate CA is not needed. Users request certificates with a certificate signing request (CSR). Wildcard certificates reduce the administrative burden for certificates, but do not have anything to do with validating certificates.

19. **A.** Online Certificate Status Protocol (OCSP) validates trust with certificates. Clients send the serial number of the certificate to the Certificate Authority (CA) within the Public Key Infrastructure (PKI) and the CA returns short responses such as good, unknown, or revoked. A certificate revocation list (CRL) includes a list of revoked certificates listed by serial numbers and can become quite large after a while. The CA isn't a protocol. You request certificates with a certificate signing request (CSR).

20. **C.** Recovery agents can decrypt data and messages if the user's private key is no longer available. Although certificate authorities use trust models, a trust relationship doesn't directly apply here. A user's public key is already publicly available, so it isn't useful here. A certificate revocation list (CRL) is a list of revoked certificates and doesn't apply in this scenario.

Chapter 11

Exploring Operational Security

CompTIA Security+ objectives covered in this chapter:

2.1 Explain the importance of risk related concepts.

- Importance of policies in reducing risk (Privacy policy, Acceptable use, Security policy, Mandatory vacations, Job rotation, Separation of duties, Least privilege)

2.2 Summarize the security implications of integrating systems and data with third parties.

- On-boarding/off-boarding business partners, Social media networks and/or applications, Interoperability agreements (SLA, BPA, MOU, ISA), Privacy considerations, Risk awareness, Unauthorized data sharing, Data ownership, Data backups, Follow security policy and procedures, Review agreement requirements to verify compliance and performance standards

2.3 Given a scenario, implement appropriate risk mitigation strategies.

- Change management, Incident management, User rights and permissions reviews, Enforce policies and procedures to prevent data loss or theft

2.4 Given a scenario, implement basic forensic procedures.

- Order of volatility, Capture system image, Network traffic and logs, Capture video, Record time offset, Take hashes, Screenshots, Witnesses, Track man hours and expense, Chain of custody, Big Data analysis

2.5 Summarize common incident response procedures.

- Preparation, Incident identification, Escalation and notification, Mitigation steps, Lessons learned, Reporting, Recovery/reconstitution procedures, First responder, Incident isolation (Quarantine, Device removal), Data breach, Damage and loss control

2.6 Explain the importance of security related awareness and training.

- Security policy training and procedures, Role-based training, Personally identifiable information, Information classification (High, Medium, Low, Confidential, Private, Public)
- Data labeling, handling, and disposal
- Compliance with laws, best practices, and standards
- User habits (Data handling, Clean desk policies, Personally owned devices), Use of social networking and P2P, Follow up and gather training metrics to validate compliance and security posture

4.4 Implement the appropriate controls to ensure data security.

- Data policies (Wiping, Disposing, Retention, Storage)

5.2 Given a scenario, select the appropriate authentication, authorization, or access control.

- Authorization (Least privilege, Separation of duties)

5.3 Install and configure security controls when performing account management, based on best practices.

- Mitigate issues associated with users with multiple account/roles and/or shared accounts

**

Organizations often develop security policies. These provide guiding principles to the professionals who implement security throughout the organization. These policies, combined with training for personnel to raise overall security awareness, help reduce security incidents. However, security incidents still occur, and incident response policies provide the direction on how to handle them.

Exploring Security Policies

Security policies are written documents that lay out a security plan within a company. They are one of many management controls used to reduce and manage risk. When created early enough, they help ensure that personnel consider and implement security throughout the life cycle of various systems in the company. When the policies and procedures are enforced, they help prevent incidents, data loss, and theft.

Policies include brief, high-level statements that identify goals based on an organization's overall beliefs and principles. After creating the policy, the organization creates guidelines and procedures to support the policies. Although the policies are often high-level statements, the guidelines and procedures provide details on policy implementation.

Security controls enforce the requirements of a security policy. For example, a security policy may state that internal users must not use peer-to-peer (P2P) applications. A firewall with appropriate rules to block these applications provides a technical implementation of this policy. Similarly, administrators can use port-scanning tools to detect the applications running on internal systems and violating the security policy.

A security policy can be a single large document or divided into several smaller documents, depending on the needs of the company. The following sections identify many of the common elements of a security policy.

Remember this

Written security policies are management controls that identify a security plan. Other security controls, such as technical, operational, and additional management controls, enforce security policies.

Personnel Policies

Companies frequently develop policies to specifically define and clarify issues related to personnel. This includes personnel behavior, expectations, and possible consequences. Personnel learn these policies when they are hired and as changes occur. Some of the policies directly related to personnel are acceptable use, mandatory vacations, separation of duties, job rotation, and clean desk policies. The following sections cover these in more depth.

Acceptable Use Policy and Privacy Policy

An acceptable use policy (AUP) defines proper system usage. It often describes the purpose of computer systems and networks, how users can access them, and the responsibilities of users when accessing the systems. Many organizations monitor user activities, such as what web sites they visit, and data they send out via email. The AUP typically includes statements informing users that systems are in place monitoring their activities. For example, a proxy server logs all web sites that a user visits.

In some cases, the AUP might include privacy statements informing users what computer activities they can consider private. Many users have an expectation of privacy when using an organization's computer systems and network that isn't justified. The privacy policy statement helps to clarify the organization's stance.

The AUP often includes definitions and examples of unacceptable use. For example, it may prohibit employees from using company resources to access peer-to-peer (P2P) sites or social media sites.

Many organizations require users to read and sign a document indicating they understand the acceptable use policy when they're hired and in conjunction with annual security training. In many cases, organizations post the policy on an intranet site and sign it electronically. Other methods, such as logon banners or emails, help reinforce an acceptable use policy.

Mandatory Vacations

Mandatory vacation policies help detect when employees are involved in malicious activity, such as fraud or embezzlement. As an example, employees in positions of fiscal trust, such as stock traders or bank employees, are often required to take an annual vacation of at least five consecutive workdays.

For embezzlement actions of any substantial size to succeed, an employee would need to be constantly present in order to manipulate records and respond to different inquiries. On the other hand, if an employee is forced to be absent for at least five consecutive workdays, the likelihood of any illegal actions succeeding is reduced because someone else would be required to answer the queries during the employee's absence.

Mandatory vacations aren't limited to only financial institutions, though. Many organizations require similar policies for administrators. For example, an administrator may be the only person required to perform sensitive activities such as reviewing logs. A malicious administrator can overlook or cover up certain activities revealed in the logs. However, a mandatory vacation would require someone else to perform these activities, which increases the chance of discovery.

Of course, mandatory vacations by themselves won't prevent fraud. Most companies will implement the principle of defense in depth by using multiple layers of protection. Additional policies may include separation of duties and job rotation to provide as much protection as possible.

Remember this

Mandatory vacation policies require employees to take time away from their job. These policies help to deter fraud and discover malicious activities while the employee is away.

Separation of Duties

Separation of duties is a principle that prevents any single person or entity from being able to complete all the functions of a critical or sensitive process. It's designed to prevent fraud, theft, and errors.

Accounting provides the classic example. It's common to divide Accounting departments into two divisions: Accounts Receivable and Accounts Payable. Personnel in the Accounts Receivable division review and validate bills. They then send the validated bills to the personnel in the Accounts Payable division, who pay the bills. Similarly, this policy would ensure personnel are not authorized to print and sign checks. Instead, a separation of duties policy separates these two functions to reduce the possibility of fraud.

If Homer were the only person doing all these functions, it would be possible for him to create and approve a bill from Homer's Most Excellent Retirement Account. After approving the bill, Homer would then pay it. If Homer doesn't go to jail, he may indeed retire early at the expense of the financial health of the company.

Separation of duties policies also apply to information technology (IT) personnel. For example, it's common to separate application development tasks with application deployment tasks. In other words, developers create and modify applications and then pass the compiled code to administrators. Administrators then deploy the code to live production systems. Without this policy in place, developers might be able make quick, untested changes to

code resulting in unintended outages. This provides a high level of version control and prevents potential issues created through uncontrolled changes.

As another example, a group of IT administrators may be assigned responsibility for maintaining a group of database servers, but do not have access to security logs on these servers. Instead, security administrators regularly review these logs, but these security administrators will not have access to data within the databases.

Consider what should happen if one of the IT administrators is promoted and is now working as a security administrator. Based on separation of duties, this administrator should now have access to security logs, but access to the data within the databases should be revoked. However, if the administrator's permissions to the data are not revoked, the administrator will have more permissions than needed, violating the principle of least privilege. A user rights and permissions review often discovers these types of issues.

Remember this

Separation of duties prevents any single person or entity from being able to complete all the functions of a critical or sensitive process by dividing the tasks between employees. This helps prevent fraud that can occur if a single person prints and signs checks.

Job Rotation

Job rotation is a concept that has employees rotate through different jobs to learn the procedures and processes in each. From a security perspective, job rotation helps to prevent or expose dangerous shortcuts or even fraudulent activity.

For example, your company could have an Accounting department. As mentioned in the "Separation of Duties" section, you would separate accounting into two divisions—Accounts Receivable and Accounts Payable. Additionally, you could rotate personnel in and out of jobs in the two divisions. This would ensure more oversight over past transactions and help ensure that employees are following rules and policies.

In contrast, imagine a single person always performs the same function without any expectation of oversight. This increases the temptation to go outside the bounds of established policies.

Job rotation policies work well together with separation of duties policies. A separation of duties policy helps prevent a single person from controlling too much. However, if an organization only used a separation of duties policy, it is possible for two people to collude in a scheme to defraud the company. If a job rotation policy is also used, these two people will not be able to continue the fraudulent activity indefinitely.

Job rotation policies also apply to IT personnel. For example, the policy can require administrators to swap roles on a regular basis, such as annually or quarterly. This prevents any single administrator from having too much control over a system or network.

Remember this

Job rotation policies require employees to change roles on a regular basis. Employees might change roles temporarily, such as for three to four weeks, or permanently. This helps ensure that employees cannot continue with fraudulent activity indefinitely.

I'll Go to Jail Before I Give You the Passwords!

The city of San Francisco had an extreme example of the dangers of a single person with too much explicit knowledge or power. A network administrator with Cisco's highest certification of Cisco Certified Internetwork Expert (CCIE) made changes to the city's network, changing passwords so that only he knew them and ensuring that he was the only person with administrative access.

It could be that he was taking these actions to protect the network that he considered his "baby." He was the only CCIE, and it's possible he thought others did not have the necessary knowledge to maintain the network adequately. Over the years, fewer and fewer people had access to what he was doing, and his knowledge became more and more proprietary. Instead of being malicious in nature, he may have simply been protective, even if overly protective.

At some point, his supervisor recognized that all the information eggs were in the basket of this lone CCIE. It was just too risky. What if a bus, or one of San Francisco's famous trolleys, hit him? What would the organization do? His supervisor asked him for some passwords and he refused, even when faced with arrest. Later, he gave law enforcement personnel passwords that didn't work.

Law enforcement personnel charged him with four counts of tampering with a computer network and courts kept him in custody with a \$5 million bail. Ultimately, a court convicted him of one felony count and sentenced him to four years in prison. This is a far fall from his reported annual salary of \$127,735.

The city of San Francisco had to bring in experts from Cisco and the city reported costs of \$900,000 to regain control of their network. Following his conviction, the court also ordered the administrator to pay \$1.5 million in restitution.

What's the lesson here? Internal security controls, such as creating and enforcing policies related to rotation of duties, separation of duties, and cross training, may have been able to avoid this situation completely. If this CCIE truly did have good intentions toward what he perceived as his network, these internal controls might have prevented him from going over the line into overprotection and looking at the world through the bars of a jail cell.

...

Clean Desk Policy

A clean desk policy directs users to keep their areas organized and free of papers. The primary security goal is to reduce threats of security incidents by ensuring the protection of sensitive data. More specifically, it helps prevent the possibility of data theft or inadvertent disclosure of information.

Imagine an attacker going into a bank for a bank loan and meeting a loan officer. The loan officer has stacks of paper on his or her desk, including loan applications from various customers. If the loan officer steps out, the attacker can easily grab some of the documents, or simply take pictures of the documents with a mobile phone.

Beyond security, organizations want to present a positive image to customers and clients. Employees with cluttered desks with piles of paper can easily turn off customers.

However, a clean desk policy doesn't just apply to employees who meet and greet customers. It also applies to employees who don't interact with customers. Just as dumpster divers can sort through trash to gain valuable information, anyone can sort through papers on a desk to learn information. It's best to secure all papers to keep them away from prying eyes. Some items left on a desk that can present risks include:

- Keys
- Cell phones
- Access cards
- Sensitive papers
- Logged-on computer
- Printouts left in printer
- Passwords on Post-it notes
- File cabinets left open or unlocked
- Personal items such as mail with Personally Identifiable Information (PII)

Some people want to take a clean desk policy a step further by scrubbing and sanitizing desks with antibacterial cleaners and disinfectants on a daily basis. They are free to do so, but that isn't part of a security-related clean desk policy.

Remember this

A clean desk policy requires users to organize their areas to reduce the risk of possible data theft. It reminds users to secure sensitive data and may include a statement about not writing down passwords.

Account Management Policies

Chapter 2, “Exploring Control Types and Methods,” covers account management as a logical access control. Accounts provide access to systems and networks, and account management involves the creation, deletion, and disabling of accounts. Account management policies provide direction for administrators to address and prevent these vulnerabilities. Two key elements stressed in Chapter 2 are:

- **Least privilege policy.** This policy ensures that users have only the rights and permissions they need for the job and no more. Administrators should follow this policy when creating and maintaining accounts. User rights and permissions reviews (described in Chapter 8, “Managing Risk”) ensure administrators are following the policy.
- **Account disablement policy.** This policy mandates that administrators disable user accounts as soon as possible when a user leaves the organization. Audits and reviews can verify if the administrators are following the policy.

The following sections discuss additional elements often included in an account management policy, along with an explanation of their value.

Require Administrators to Use Two Accounts

It’s common to require administrators to have two accounts. They use one account for regular day-to-day work. It has the same limited privileges as a regular user. The other account has elevated privileges required to perform administrative work, and they use this only when performing administrative work. The benefit of this practice is that it reduces the exposure of the administrative account to an attack.

For example, when malware infects a system, it often attempts to gain additional rights and permissions using privilege escalation techniques. It may exploit a bug or flaw in an application or operating system. Or, it may simply assume the rights and permissions of the logged-on user.

If an administrator logs on with an administrative account, the malware can assume these elevated privileges. In contrast, if the administrator is logged on with a regular user account, the malware isn’t able to escalate its privileges through this account.

This also reduces the risk to the administrative account for day-to-day work. Imagine Homer is an administrator and he’s called away to a crisis. It is very possible for him to walk away without locking his computer. If he was logged on with his administrator account, an attacker walking by can access the system and have administrative privileges. Although systems often have password-protected screen savers, these usually don’t start until about 10 minutes or longer after a user walks away.

Never Use Shared Accounts

Account management policies often dictate that personnel should not use shared accounts. Instead, each user has at least one account, which is only accessible to that user. If multiple users share a single account, you aren’t able to implement basic authorization controls. Chapter 1, “Mastering Security Basics,” discusses authentication concepts in depth. As a reminder, three key concepts are:

- **Identification.** Users claim an identity with an identifier such as a username.
- **Authentication.** Users provide their identity using an authentication method such as a password.
- **Authorization.** Users are authorized access to resources, based on their proven identity.

Imagine that Bart, Maggie, and Lisa all used a Guest account. If you want to give Lisa access to certain files, you’d grant access to the Guest account, but Bart and Maggie would have the same access. If Bart deleted the files, logs would indicate the Guest account deleted the files, but you wouldn’t know who actually deleted the files. In contrast, if users have unique user accounts, you can give them access to resources individually. Additionally, logs would indicate exactly who took an action.

Remember this

Requiring administrators to use two accounts, one with administrator privileges and another with regular user privileges, helps prevent privilege escalation attacks. Users should not use shared accounts.

Third-Party Issues

There are multiple instances where an organization works with another organization as a third party and it can bring up a variety of security issues. A third party is an entity that isn't directly involved in activities between two primary parties.

As an example, imagine that an organization called Get Certified Get Ahead hosts a web site (<http://getcertifiedgetahead.com/>) and collects customer data. The *GetCertifiedGetAhead.com* web site and each customer make up the primary two parties. Now, imagine that Get Certified Get Ahead hires an application developer to modify the web site application, and the developer can potentially access user data. In this situation, the application developer is the third party and represents a potential risk when granted access to the web site and data.

On-boarding business partners refers to granting them access to a system by giving them an account. One of the key considerations when on-boarding a business partner is to ensure you follow the principle of least privilege. Grant them access to what they need for their job, but no more. Off-boarding is the process of removing their access. Once a partnership has ended, you need to revoke the business partner's access.

In many situations, it's appropriate to use a non-disclosure agreement (NDA) to ensure that third parties understand their responsibilities. This can be a completely separate agreement, but is more commonly embedded as a clause in a contract with the third party. Some concerns or issues to include in the contract or NDA are:

- **Privacy considerations.** A contract stresses the importance of maintaining privacy of any data, and especially PII. The goal is to ensure that the third party understands the importance of protecting the data.
- **Data ownership.** The first party owns the data and a contract might stress this. In other words, the third party is not authorized to keep, use, or distribute the data and should destroy it upon completion of the contract period.
- **Data backups.** If activity by the third party has the potential to destroy or corrupt data, a contract might stress the need for the third party to maintain up-to-date backups.
- **Unauthorized data sharing.** Because the third party does not own the data, it is not authorized to share it with any other entities.
- **Security policy and procedures.** During the period of the contract, the third party is responsible for following appropriate security policies and procedures. This might include items such as encrypting all data in transit or reporting potential incidents as soon as they are discovered.
- **Reviews.** Depending on the scope of the partnership, an NDA might include review clauses. The goal of a review is to verify the third party is complying with the agreement and meeting performance standards.

Remember this

When working with third parties or as a third party, it's important to protect data. Most non-disclosure agreements prohibit sharing data unless you are the data owner.

Interoperability Agreements

In addition to NDAs, organizations often utilize different interoperability agreements to identify various responsibilities. These include:

- **Interconnection security agreement (ISA).** An ISA specifies technical and security requirements for planning, establishing, maintaining, and disconnecting a secure connection between two or more entities. For example, it may stipulate certain types of encryption for all data in transit.
- **Service level agreement (SLA).** An SLA is an agreement between a company and a vendor that stipulates performance expectations, such as minimum uptime and maximum downtime levels. Organizations use SLAs when contracting services from service providers such as Internet Service Providers (ISPs). Many SLAs include a monetary penalty if the vendor is unable to meet the agreed-upon expectations.
- **Memorandum of understanding (MOU).** An MOU expresses an understanding between two or more parties indicating their intention to work together toward a common goal. It is similar to an SLA in that it defines the responsibilities of each of the parties. However, it is less formal than an SLA and does not include monetary penalties. Additionally, it doesn't have strict guidelines in place to protect sensitive data. Many times, MOUs are used in conjunction with ISAs. National Institute of Standards and Technology (NIST) Special Publication (SP) 800-47, "Security Guide for Interconnecting Information Technology Systems," includes more in-depth information on MOUs and ISAs.
- **Business partners agreement (BPA).** A BPA is a written agreement that details the relationship between business partners, including their obligations toward the partnership. It typically identifies the shares of profits or losses each partner will take, their responsibilities to each other, and what to do if a partner chooses to leave the partnership. One of the primary benefits of a BPA is that it can help settle conflicts when they arise.

Remember this

A memorandum of understanding (MOU) defines responsibilities of each party, but it is not as strict as a service level agreement (SLA) or interconnection security agreement (ISA). If the parties will be handling sensitive data, they should include an ISA to ensure strict guidelines are in place to protect the data while in transit.

Change Management Policy

The worst enemies of many networks have been unrestrained administrators. A well-meaning administrator can make what appears to be a minor change to fix one problem, only to cause a major problem somewhere else. A misconfiguration can take down a server, disable a network, stop email communications, and even stop all network traffic for an entire enterprise.

For example, I once saw a major outage occur when an administrator was troubleshooting a printer problem. After modifying the printer's Internet Protocol (IP) address, the printer began to work. Sounds like a success, doesn't it? Unfortunately, the new IP address was the same IP address assigned to a Domain Name System (DNS) server, and it created an IP address conflict. The conflict prevented the DNS server from resolving names to IP addresses. This resulted in a major network outage until another administrator discovered and corrected the problem.

These self-inflicted disasters were relatively common in the early days of IT. They still occur today, but organizations with mature change management processes in place have fewer of these problems. Change management defines the process for any type of system modifications or upgrades. It provides two key goals:

- To ensure changes to IT systems do not result in unintended outages
- To provide an accounting structure or method to document all changes

When a change management program is in place, administrators are discouraged from making configuration changes without submitting the change for review and approval. In other words, they don't immediately make a change as soon as they identify a potential need for the change. This includes making any type of configuration changes to systems, applications, patches, or any other change. Instead, they follow the change management process before making a change.

Experts from different areas of an organization examine change requests and can either approve or postpone them. The process usually approves simple changes quickly. A formal change review board regularly reviews postponed requests and can approve, modify, or reject the change.

This entire process provides documentation for approved changes. For example, some automated change management systems create accounting logs for all change requests. The system tracks the request from its beginning until implementation. Administrators use this documentation for configuration management and disaster recovery. If a modified system fails, change and configuration management documentation identifies how to return the system to its prefailure state.

Chapter 5, "Securing Hosts and Data," covers patch management policies. Patch management ensures systems are kept up to date and reduces risks associated with known vulnerabilities. However, patches can cause unintended outages, so many organizations include patch management processes within a change management process. When patch management is included with change management, the change management process provides documentation of the patches.

Remember this

Change management defines the process and accounting structure for handling modifications and upgrades. The goals are to reduce risks related to unintended outages and provide documentation for all changes.

Data Policies

Every company has secrets. Keeping these secrets can often make the difference between success and failure. A company can have valuable research and development data, customer databases, proprietary information on products, and much more. If the company cannot keep private and proprietary data secret, it can directly affect its bottom line.

Data policies assist in the protection of data and help prevent data leakage. This section covers many of the different elements that may be contained in a data policy.

Information Classification

As a best practice, organizations take the time to identify, classify, and label data they use. Data classifications ensure that users understand the value of data, and the classifications help protect sensitive data. Classifications can apply to hard data (printouts) and soft data (files).

As an example, the U.S. government uses classifications such as *Top Secret*, *Secret*, *Confidential*, and *Unclassified* to identify the sensitivity of data. Private companies often use terms such as *Proprietary*, *Private*, *Classified*, or *Public*. Some companies simply use the terms *high*, *medium*, and *low* to describe the value of the data (as in high value, medium value, and low value).

The labels and classifications an organization uses are not as important as the fact that it uses labels and classifications. Organizations take time to analyze their data, classify it, and provide training to users to ensure the users recognize the value of the data. They also include these classifications within a data policy.

Data Labeling and Handling

Data labeling ensures that users know what data they are handling and processing. For example, if an organization classified data as confidential, private, sensitive, and public, it would also use labeling to identify the data. These labels can be printed labels for media such as backup tapes. It's also possible to label files using file properties, headers, footers, and watermarks.

Consider a company that spends millions of dollars on research and development (R&D) trying to develop or improve products. The company values this data much more than data publicly available on its web site, and it needs to protect it. However, if employees have access to the R&D data and it's not classified or labeled, they may not realize its value and may not protect it.

For example, a web content author may write an article for the company's web site touting its achievements. If the R&D data isn't classified and labeled, the author may include some of this R&D data in the article, inadvertently giving the company's competitors free access to valuable data. Although the R&D employees will easily recognize the data's value, it's not safe to assume that everyone does. On the other hand, if the data includes confidential or proprietary labels, anyone would recognize its value and take appropriate steps to protect it.

Chapter 9, "Preparing for Business Continuity," presented information on backups. As a reminder, it's important to protect backups with the same level of protection as the original data. Labels on backup media help administrators easily identify the value of the data on the backups.

Remember this

Data classifications and data labeling help ensure personnel apply the proper security controls to protect information.

Data Wiping and Disposing

When computers reach the end of their life cycles, organizations donate them, recycle them, or sometimes just throw them away. From a security perspective, you need to ensure that the computers don't include any data that may be useful to people outside your organization or damaging to your organization if unauthorized people receive it.

It's common for organizations to have a checklist to ensure that personnel sanitize a system prior to disposing of it. The goal is to ensure that personnel remove all usable data from the system.

Hard drives represent the greatest risk because they hold the most information, so it's important to take additional steps when decommissioning old hard drives. Simply deleting a file on a drive doesn't actually delete it. Instead, it marks the file for deletion and makes the space available for use. Similarly, formatting a disk drive won't erase the data. There are many recovery applications available to recover deleted data, file remnants, and data from formatted drives.

Instead, technicians use different methods to wipe all the data off drives before disposing of them. These methods sanitize the drives, ensuring that they do not contain any valuable information. Some methods used to sanitize drives are:

- **Bit-level overwrite.** Different programs are available that write patterns of 1s and 0s multiple times to ensure that data originally on the disk is unreadable. This process ensures that the disk doesn't contain any data.
- **Degauss the disks.** A degausser is a very powerful electronic magnet. Passing a disk through a degaussing field renders the data on the disk unreadable, and it often destroys the motors of the disk. Degaussing of backup tapes sanitizes a tape without destroying it.

- **Physical destruction.** If the disk includes classified or proprietary data, simply overwriting it may not be enough. Instead, the computer disposal policy may require the destruction of the drive. For example, technicians can remove disk platters and sand them down to the bare metal.

It's also worth mentioning that hard drives can be in other devices besides computers. For example, many copy machines include disk drives, and they can store files of anything that employees recently copied or printed. If personnel don't sanitize the disk drives, it can also result in a loss of confidentiality.

Similarly, organizations often have a policy related to paper containing any type of proprietary or private data. Shredding or incinerating these papers prevents them from falling into the wrong hands. If personnel just throw this paper away, dumpster divers can sift through the trash and gain valuable information.

Wiping Files

In some cases, technicians want to erase specific files and ensure that a system doesn't have any remnants of these files on it. Erasing a single file is typically done by overwriting it similar to a bit-level overwrite for an entire drive. Many antivirus programs have shredding tools, which ensure the file is no longer accessible.

There are some situations where you want to keep the files, but ensure that none of the files holds random data. As an example, imagine that someone inadvertently began working with a proprietary or secret file on a system that should only hold public or unclassified data. You can shred the file, but it's still possible that the system holds remnants of data from the classified file. Cluster tip wiping tools help erase all remnants of the data. More specifically, these tools erase remnants contained at the end or tip of the last cluster of a file.

Files are stored in clusters and cluster sizes are typically about 4 KB. Files use as many clusters as they need, but the last cluster has some unused space that the operating system pads with random data. For example, imagine you are saving a 6 KB file. It will use two 4 KB clusters and the last 2 KB in the second cluster isn't used to store information for your file. However, this last 2 KB isn't empty. Instead, it contains random data pulled from memory. If someone was recently working with proprietary or secret data, the last 2 KB might hold some of that data. Cluster tip wiping tools can sanitize files stored on a system, and eliminate this issue.

Storage and Retention Policies

A storage and retention policy identifies where data is stored and how long it is retained. For example, a storage policy often dictates that users must store all data on servers instead of local workstations. One of the benefits is that administrators can back up data on the server to ensure they have copies of user data. If users store data on their individual systems, it makes it much more difficult and expensive to back up data.

Retention policies help reduce legal liabilities, and this is another reason they're used. For example, imagine if a retention policy states that the company will only keep email for one year. A court order requiring all email from the company can only expect to receive email from the last year.

On the other hand, if the organization doesn't have a retention policy, it may need to provide email from the past 10 years or longer in response to a court order. This can require an extensive amount of work by administrators to recover archives or search for specific emails. Additionally, investigations can uncover other embarrassing evidence from previous years. The retention policy helps avoid these problems.

Some laws mandate the retention of data for specific time frames, such as three years or longer. For example, laws mandate the retention of all White House emails indefinitely. If a law applies to an organization, the retention policy reflects the same requirements.

Personally Identifiable Information

Personally Identifiable Information (PII) is personal information that can be used to personally identify an individual. Some examples of PII are:

- Full name
- Birthday and birth place
- Medical and health information
- Street or email address information
- Personal characteristics, such as biometric data
- Any type of identification number, such as a Social Security number (SSN) or driver's license number

In general, you need two or more pieces of information to make it PII. For example, "John Smith" is not PII by itself because it can't be traced back to a specific person. However, when you connect the name with a birth date, an address, medical information, or other data, it is PII.

When attackers gain PII, they often use it for financial gain at the expense of the individual. For example, attackers steal identities, access credit cards, and empty bank accounts. Whenever possible, organizations should minimize the use, collection, and retention of PII. If it's not kept, it can't be compromised. On the other hand, if a company collects PII and attackers compromise the data, the company is liable.

The number of security breach incidents resulting in the loss of PII continues to rise. For example, a Veterans Affairs (VA) employee copied a database onto his laptop that contained PII on over 26 million U.S. veterans. He took the laptop home and a burglar stole it. The VA then went through the

painful and expensive process of notifying all of the people who were vulnerable to identity theft, and the affected individuals spent countless hours scouring their records for identity theft incidents. Even though police later recovered the laptop, the VA paid \$20 million to settle a lawsuit in the case.

Chapter 5 mentioned several other instances, such as the attack on Sony’s PlayStation Network that compromised more than 77 million customer records, resulting in direct expense of over \$171 million.

Each of these instances resulted in potential identity theft and the loss of goodwill and public trust of the company. Both customers and employees were negatively impacted, and the companies were forced to spend time and energy discussing the incident, and spend money trying to repair their reputations.

Protecting PII

Organizations have an obligation to protect PII. There are many laws that mandate the protection of PII, including international laws, federal laws, and local regulations. Organizations often develop policies to identify how they handle, retain, and distribute PII, and these policies help ensure they are complying with relevant regulations. When a company doesn’t use a specific PII policy, it usually identifies methods used to protect PII in related data policies.

Many laws require a company to report data losses due to security breaches. If an attack results in the loss of customer PII data, the company is required to report it and notify affected individuals. As an example, Arizona enacted a security breach notification law that requires any company doing business in Arizona to notify customers of security breaches. Most states in the United States have similar laws, and similar international laws exist.

One of the common reasons data seems to fall into the wrong hands is that employees don’t understand the risks involved. They may not realize the value of the data on a laptop, or they may casually copy PII data onto a USB flash drive. As mentioned previously, data classification and labeling procedures help employees recognize the data’s value, and help protect sensitive data.

Training is also important. One of the goals of security professionals is to reinforce the risks of not protecting PII. When employees understand the risks, they are less likely to risk customer and employee data to identity theft.

Additionally, if employees need to transmit PII over a network, they can ensure it’s protected by using encryption. As mentioned previously in this book, encrypting data in transit provides strong protection against loss of confidentiality.

Remember this

Personally Identifiable Information (PII) includes information such as a full name, birth date, biometric data, and identifying numbers such as a SSN. Organizations have an obligation to protect PII and often identify procedures for handling and retaining PII in data policies.

Privacy Policy

It’s almost a business requirement today for a company to have a web site. Customers expect a web site and often look for it to get additional information about a company. When it doesn’t exist, customers often go elsewhere. However, web sites have additional requirements such as a privacy policy.

A privacy policy identifies how a web site collects, uses, and discloses information about visitors. For example, web forms collect email addresses and other information from users. The privacy policy indicates whether the company uses this information internally only or if it sells or shares it with other entities.

Many states, such as California, Nebraska, and Pennsylvania, have specific laws requiring privacy policies. For example, a California law requires web sites to post a privacy policy in plain view on the site. This law applies to any web site that collects information about California residents, regardless of where the web site is located.

You can usually find a link to a privacy policy on the site’s main page. For example, if you go to Google.com, you’ll find a link labeled “Privacy,” and by clicking on it, you’ll see its privacy policy.

Social Media Networks and Applications

Millions of people interact with each other using social media networks and applications such as Facebook and Twitter. Facebook allows people to share their lives with friends, family, and others. Twitter allows people to tweet about events as they are happening. From a social perspective, these technologies allow people to share information about themselves with others. A user posts a comment and a wide group of people instantly sees it.

However, from a security perspective, they present some significant risks, especially related to inadvertent information disclosure. Attackers can use these sites to gain information about individuals and then use that information in an attack. Organizations typically either train users about the risks or block access to the social media sites to avoid the risks.

Users often post personal information, such as birth dates, their favorite colors or books, the high school they graduated from, graduation dates, and much more. Some sites use this personal information to validate users when they forget or need to change their password. For example, imagine Maggie needs to reset her password for a bank account. The web site may challenge her to enter her birth date, favorite book, and graduation date for validation.

This is also known as a cognitive password and, theoretically, only Maggie knows this information. However, if Maggie posts all this information on Facebook, an attacker can use it to change the password on the bank account.

As an example, David Kernell used Yahoo!’s cognitive password account recovery process to change Governor Sarah Palin’s password for her email account. At the time, Yahoo! asked questions such as her high school and birth date and Kernell obtained all the information from online searches. Of course, it didn’t turn out well for him. A jury convicted him of a felony and he served more than a year in prison.

In some cases, attackers have used personal information from social networking sites to launch scams. For example, attackers first identify the name of a friend or relative using the social networking site. The attackers then impersonate the friend or relative in an email, claiming they were robbed and are stuck in a foreign country. Attackers end the email with a plea for help asking the victim to send money via wire transfer.

It’s also worth noting that social networking sites have become one of the methods that employers use to collect information on prospective employees. In 2010, Microsoft surveyed U.S. human resources professionals and learned that 70 percent of them had rejected a job application based on information they found online.

Remember this

Social media sites allow people to share personal comments with a wide group of people. However, improper use of social networking sites can result in inadvertent information disclosure. Attackers can also use information available on these sites to launch attacks against users or in a cognitive password attack to change a user’s password. Training helps users understand the risks.

SSO and Social Media

Social media sites have implemented single sign-on (SSO) capabilities that are easy to integrate into other web sites. For example, if you log on to Facebook, you can then access other web sites without logging on again. Similarly, if you have your Facebook passwords stored in your web browser, and you visit another web site, you can click on a button to log on with your Facebook password.

Although this is convenient, it also poses additional risks. For example, a data breach on the social media site can expose the user passwords to attackers and attackers can then use these passwords to access additional applications or web sites. This becomes a balance between risk and convenience. The risk is minimal as long as an attacker cannot access any valuable data such as financial data. However, if the attacker can access any type of valuable data, it’s best to avoid linking the accounts together.

Additionally, if users have the same password on multiple sites, a data breach on any site can give attackers the information they need to log on to other sites. It’s best to use different passwords for different sites. At the very least, use different passwords for sites that provide access to valuable data.

Banner Ads and Malvertisements

Attackers have been delivering malware through malicious banner ads for several years now. These look like regular ads, but they contain malicious code. Many of these are Flash applets with malicious code embedded in them, but others just use code to redirect users to another server, such as one with a drive-by download waiting for anyone who clicks.

Although these malvertisements have been on many social media sites, they’ve also appeared on mainstream sites. For example, attackers installed a malvertisement on the *New York Times* web site where it ran for about 24 hours before webmasters discovered and disabled it.

Similarly, malvertising has appeared on the Yahoo! web site. In late 2013 and early 2014, users who clicked on some Yahoo! ads were taken to sites hosting fake antivirus software. These sites included pop-ups indicating that users’ systems were infected with malware and encouraging the users to download and install free antivirus software to remedy the infection. However, users who took the bait installed malware onto their systems. Later in July and August 2014, some ads on Yahoo! sent users to sites in Eastern Europe that were hosting CryptoWall, according to research by Blue Coat Systems Inc. CryptoWall is a malicious form of ransomware that encrypts user files and demands payment to decrypt them.

Attackers have used two primary methods to get these malvertisements installed on legitimate web sites. One method is to attack a web site and insert ads onto that web site. The second method is to buy ads. They often represent an ad agency pretending to represent legitimate clients. For example, one attacker convinced Gawker Media to run a series of Suzuki advertisements, which were actually malvertisements. Similarly, it’s unlikely that Yahoo! was aware that it was hosting malvertising, but instead, these ads likely appeared as a result of attacks or by being tricked.

Social Networking and P2P

Peer-to-peer (P2P or file sharing) applications allow users to share files such as music, video, and data over the Internet. Instead of a single server providing the data to end users, all computers in the P2P network are peers, and any computer can act as a server to other clients.

The first widely used P2P network was Napster, an online music-sharing service that operated between 1999 and 2001. Users copied and distributed MP3 music files among each other, and these were often pirated music files. The files were stored on each user’s system, and as long as the system was accessible on the Internet, other users could access and download the files. A court order shut down Napster due to copyright issues, but it later reopened as an online music store. Other P2P software and P2P networks continue to appear and evolve.

Organizations usually restrict the use of P2P applications in networks, but this isn't because of piracy issues. One reason is because the P2P applications can consume network bandwidth, slowing down other systems on the network. Worse, a significant risk with P2P applications is data leakage. Users are often unaware of what data they are sharing. Another risk is that users are often unaware of what data the application downloads and stores on their systems, causing them to host inappropriate data. Two examples help illustrate these data leakage risks.

Information concentrators search P2P networks for information of interest and collect it. In March 2009, investigators discovered an information concentrator in Iran with over 200 documents containing classified and secret U.S. government data. This included classified information about Marine One, the helicopter used by the president. Although the information about Marine One made the headlines, the attackers had much more information. For example, this concentrator included Iraq status reports and lists of soldiers with privacy data.

How did this happen? Investigations revealed that a defense contractor installed a P2P application on a computer. The computer had access to this data, and the P2P application shared it.

The media latched onto the news about Marine One, so this story was widely published. However, it's widely believed that much more data is being mined via P2P networks. Most end users don't have classified data on their systems, but they do have PII, such as banking information or tax data. When an attacker retrieves data on a user's system and empties a bank account, it may be a catastrophe to the user, but it isn't news.

A second example affected a school-age child. It's popular to use these P2P sharing programs to share music files, but they are often used to share other data. One school-age girl was browsing data she found on her computer and discovered a significant number of pornographic pictures. She did not seek these or deliberately download them. Instead, as a member of the P2P network, the P2P application used her system to store files shared by others.

Organizations can restrict access to P2P networks by blocking access in firewalls. Additionally, port scanners can scan open ports of remote systems to identify P2P software. Organizations often include these checks when running a port scanner as part of a vulnerability scan.

Remember this

Data leakage occurs when users install P2P software and unintentionally share files. Organizations often block P2P software at the firewall.

Responding to Incidents

Many organizations create incident response policies to help personnel identify and respond to incidents. A security *incident* is an adverse event or series of events that can negatively affect the confidentiality, integrity, or availability of data or systems within the organization, or that has the potential to do so.

Some examples include attacks, release of malware, security policy violations, unauthorized access of data, and inappropriate usage of systems. For example, an attack resulting in a data breach is an incident. Once the organization identifies a security incident, it will respond based on the incident response policy.

Organizations regularly review and update the policy. Reviews might occur on a routine schedule such as annually, or in response to an incident after performing a lessons learned review of the incident.

As an example, in the early days of computers, one hacker broke into a government system and the first thing he saw was a welcome message. He started poking around, but authorities apprehended him. Later, when the judge asked him what he was doing, he replied that when he saw the welcome message, he thought it was inviting him in. The lesson learned here was that a welcome message can prevent an organization from taking legal action against an intruder, and government systems no longer have welcome messages. Instead, they have warning banners stressing that only authorized personnel should be accessing the system. It's common to see similar warning banners when logging on to any system today.

Remember this

An incident response policy defines an incident and incident response procedures. Incident response procedures start with preparation to prepare for and prevent incidents. Preparation helps prevent incidents such as malware infections. Personnel review the policy periodically, and in response to lessons learned after incidents.

Incident Response Team

An incident response team (IRT) is composed of employees with expertise in different areas. Organizations refer to the team as a computer incident response team (CIRT), security incident response team (SIRT), or simply IRT. Combined, they have the knowledge and skills to respond to an incident. Team members may include:

- **Senior management.** Someone needs to be in charge with enough authority to get things done.
- **Network administrator/engineer.** A technical person needs to be included who can adequately understand technical problems and relay the issue to other team personnel.
- **Security expert.** Security experts know how to collect and analyze evidence using different forensic procedures.
- **Communications expert.** If an incident needs to be relayed to the public, a public relations person should be the one to do so.

Due to the complex nature of incidents, the team often has extensive training. Training includes concepts, such as how to identify and validate an incident, how to collect evidence, and how to protect the collected evidence.

Incident Response Procedures

Incident response includes multiple steps, beginning with creating the incident response policy. National Institute of Standards and Technology (NIST) Special Publication (SP) 800-61 Revision 2, “Computer Security Incident Handling Guide,” provides comprehensive guidance on how to respond to incidents. It is 79 pages so it’s obviously more in-depth than this section, but if you want to dig deeper into any topics, it’s an excellent resource.

As an overview, typical incident response procedures and concepts are:

- **Preparation.** This stage occurs before an incident and provides guidance to personnel on how to respond to an incident. It includes establishing incident response procedures and periodically reviewing and updating them. It also includes establishing procedures to prevent incidents. For example, preparation includes implementing security controls to prevent malware infections.
- **First responder.** First responders are the first security-trained individuals who arrive on the scene. The term comes from the medical community, where the first medically trained person to arrive on the scene of an emergency or accident is a first responder. A first responder could be someone from the incident response team or someone with adequate training to know what the first response steps should be. The incident response policy documents initial steps or at least the goals of first responders. In some situations, first responders might have a mini-toolkit to perform basic tests, along with a list of personnel to contact after verifying an incident occurred.
- **Incident identification.** All events aren’t security incidents so when a potential incident is reported, personnel take the time to verify it is an actual incident. For example, intrusion detection systems (IDSs) might falsely report an intrusion, but administrators would investigate it and verify it is a false positive. A false positive isn’t an actual incident. If the incident is verified, personnel might try to isolate the system based on established procedures.
- **Incident isolation.** After identifying an incident, security personnel attempt to isolate or contain it. This might include quarantining a device or removing it from the network. This can be as simple as unplugging the system’s network interface card to ensure it can’t communicate on the network. Similarly, you can isolate a network from the Internet by modifying access control lists on a router or a network firewall. This is similar to how you’d respond to water spilling from an overflowing sink. You wouldn’t start cleaning up the water until you first turn off the faucet. The goal of isolation is to prevent the problem from spreading to other areas or other computers in your network, or to simply stop the attack.
- **Damage and loss control.** When isolating an incident and throughout the entire incident response procedure, personnel attempt to limit damages and losses. Methods vary depending on the incident. As one example, the organization might identify a public relations specialist to communicate with the media, and limit potential damage from bad publicity.
- **Escalation and notification.** After identifying the incident and isolating the system, personnel escalate the incident by notifying appropriate personnel. For example, a first responder might notify an incident response team if the organization has one. If a team isn’t established, the first responder may instead identify security or forensic experts about the incident, based on established policies. The incident response policy will typically list other personnel to inform such as security managers within the organization. Forensic experts may begin a forensic evaluation depending on the scope of the incident.
- **Reporting.** In some situations, security personnel may need to notify executives within the company of the incident. Obviously, they wouldn’t notify executives of every single incident. However, they would notify executives about serious incidents that have the potential to affect critical operations. Additionally, some incidents require an organization to notify personnel outside the organization, such as customers.
- **Data breach.** If the incident involves a data breach, personnel need to identify the extent of the loss, and determine if outside entities are affected. For example, if attackers successfully attacked a system and collected customer data such as credit information, the organization has a responsibility to notify customers of the data breach as soon as possible.
- **Recovery/reconstitution procedures.** After the forensic evidence collection process, administrators will recover or restore the system to bring it back into service. Recovery or reconstitution of a system may require a simple reboot or it may require a complete rebuild of the system, depending on the incident. If the system needs to be rebuilt, it’s important to ensure that all updates and patches are also applied. Change management logs are invaluable during this process.
- **Lessons learned.** After personnel manage an incident, security personnel perform a lessons learned review. It’s very possible the incident provides some valuable lessons and the organization may modify procedures or add additional controls to prevent a reoccurrence of the incident. A review might indicate a need to provide additional training to users, or indicate a need to update the incident response policy. The goal is to prevent a future reoccurrence of the incident.
- **Mitigation steps.** After security personnel complete the review of lessons learned, they typically provide recommendations to mitigate similar risks in the future. For example, if an attack was successful because router operating systems were out of date, security personnel may update patch management policies to ensure administrators keep routers up to date.

Remember this

After identifying an incident, personnel attempt to contain or isolate the problem. This is often as simple as

disconnecting a computer from a network. Reviewing lessons learned allows personnel to analyze the incident and the response with a goal of preventing a future occurrence.

Implementing Basic Forensic Procedures

A forensic evaluation helps the organization collect and analyze data as evidence it can use in the prosecution of a crime. In general, forensic evaluations proceed with the assumption that the data collected will be used as evidence in court. Because of this, forensic practices protect evidence to prevent modification and control evidence after collecting it.

Once the incident has been contained or isolated, the next step is a forensic evaluation. What do you think of when you hear *forensics*? Many people think about the TV program *CSI* (short for “crime scene investigation”) and all of its spin-offs. These shows demonstrate the phenomenal capabilities of science in crime investigations.

Computer forensics analyzes evidence from computers to determine details on computer incidents, similar to how CSI personnel analyze evidence from crime scenes. It uses a variety of different tools to gather and analyze computer evidence. Computer forensics is a growing field, and many educational institutions offer specialized degrees around the science. Although you may not be the computer forensics expert analyzing the evidence, you should know about some of the basic concepts related to gathering and preserving the evidence.

Forensic experts use a variety of forensic procedures to collect and protect data after an attack. A key part of this process is preserving the evidence. In other words, they ensure that they don’t modify the data as they collect it, and they protect it after collection. A rookie cop wouldn’t walk through a pool of blood at a crime scene, at least not more than once. Similarly, employees shouldn’t access systems that have been attacked or power them down.

For example, files have properties that show when they were last accessed. However, in many situations, accessing the file modifies this property. This can prevent an investigation from identifying when an attacker accessed the file. Additionally, data in a system’s memory includes valuable evidence, but turning a system off deletes this data. In general, first responders do not attempt to analyze evidence until they have taken the time to collect and protect it.

Forensic experts have specialized tools they can use to capture data. For example, many experts use EnCase by Guidance Software or Forensic Toolkit by AccessData. These tools can capture data from memory or disks. This includes documents, images, email, webmail, Internet artifacts, web history, chat sessions, compressed files, backup files, and encrypted files. They can also capture data from smartphones and tablets.

Order of Volatility

Order of volatility refers to the order in which you should collect evidence. *Volatile* doesn’t mean it’s explosive, but rather that it is not permanent. In general, you should collect evidence starting with the most volatile and moving to the least volatile.

For example, random access memory (RAM) is lost after powering down a computer. Because of this, it is important to realize you shouldn’t power a computer down if it’s suspected to be involved in a security incident.

A processor can only work on data in RAM, so all the data in RAM indicates what the system was doing. This includes data users have been working on, system processes, network processes, application remnants, and much more. All of this can be valuable evidence in an investigation, but if a rookie technician turns the computer off, the evidence is lost.

Many forensic tools include the ability to capture volatile data. Once it’s captured, experts can analyze it and gain insight into what the computer and user were doing.

In contrast, data on disks remains on the drive even after powering a system down. This includes any files and even low-level data such as the Master Boot Record on a disk. However, it’s important to protect the data on the disk before analyzing it, and a common method is by capturing an image of the disk.

The order of volatility from most volatile to least volatile is:

- Data in cache memory, including the processor cache and hard drive cache
- Data in RAM, including system and network processes
- Swap file or paging file on the system disk drive
- Data stored on local disk drives
- Logs stored on remote systems
- Archive media

In case you don’t remember from your CompTIA A+ days, the swap file is an extension of RAM and it is stored on the hard drive. However, the swap file isn’t a typical file and it’s rebuilt when the system is rebooted, making it more volatile than other files stored on hard drives.

Remember this

When collecting data for a forensic analysis, you should collect it from the most volatile to the least volatile. The order of volatility is cache memory, regular RAM, swap or paging file, hard drive data, logs stored on remote systems, and archived media.

Capture System Image

A forensic image of a disk captures the entire contents of the drive. Some tools use bit-by-bit copy methods that can read the data without modifying it. Other methods include hardware devices connected to the drive to write-protect it during the copy process.

Chapter 5 introduced disk images as a common method used to deploy systems. These system disk images include mandatory security configurations and help ensure a system starts in a secure state. A distinct difference between standard system images and forensic images is that a forensic image is an exact copy and does not modify the original. This isn't always true with system imaging tools.

One of the oldest disk imaging tools used for forensics is the dd command available in Unix and Linux systems, and can be installed on Windows systems. However, other tools such as EnCase and Forensic Toolkit are much easier to use so experts don't use dd as often anymore.

All of these methods capture the entire contents of the disk, including system files, user files, and files marked for deletion but not overwritten. Similarly, many tools include the ability to capture data within volatile memory and save it as an image.

After capturing an image, experts create a copy and analyze the copy. They do not analyze the original disk and often don't even analyze the original image. They understand that by analyzing the contents of a disk directly, they can modify the contents. By creating and analyzing forensic copies, they never modify the original evidence.

Take Hashes

Hashing is an important element of forensic analysis to provide proof that collected data has retained integrity. Chapter 10, "Understanding Cryptography," covers hashes and hashing. As a reminder, a *hash* is simply a number. You can execute a hashing algorithm against data as many times as you like, and as long as the data is the same, the hash will be the same. The focus in Chapter 10 was on using hashes with files and messages. A captured forensic image (from RAM or a disk) is just a file, and you can use hashing with forensic images to ensure image integrity.

For example, after capturing an image of a disk, an expert can create a hash of the image. The expert can then write-protect the image to prevent accidental modifications during the analysis. Later, the expert can take another hash of the image and compare it with the original hash. As long as both hashes are the same, it provides proof that the image is the same and the analysis did not modify it.

Forensic analysts sometimes make a copy of the image to analyze, instead of analyzing the first image they capture. If they ever need to verify the integrity of the copy, they run the same hashing algorithm against it. Again, as long as the hash is the same, they know the analyzed data is the same as the captured data.

Similarly, some tools allow you to create a hash of an entire drive. These verify that the imaging process has not modified data. For example, you can create a hash of a drive before capturing the image and after capturing the image. If the hashes are the same, it verifies that the imaging process did not modify the drive.

Remember this

A forensic image is a bit-by-bit copy of the data and does not modify the data during the capture. Experts capture an image of the data before analysis to preserve the original and maintain its usability as evidence. Hashing provides integrity for captured images, including images of both memory and disk drives. You can take a hash of a drive before and after capturing an image to verify that the imaging process did not modify the drive contents.

Network Traffic and Logs

A forensic investigation often includes an analysis of network traffic and available logs. This information helps the investigators re-create events leading up to and during an incident.

As an example, an organization may want to prove that a specific computer was involved in an attack. One way is to match the media access control (MAC) address used by the attacking computer with an existing computer. The MAC address is permanently assigned to a network interface card, and even though the operating system can be manipulated to use a different MAC, the actual MAC isn't changed. In contrast, the IP address and name of the computer are not permanently assigned, and it is relatively easy to change them.

Chapter 8 covered protocol analyzers, and Figure 8.1 showed a capture with an expanded packet. Data within packets identifies computers involved in a conversation based on their IP address and their MAC address. If a data capture shows a MAC address matches the actual MAC address of a suspected computer, it provides a strong indication the computer was involved in the attack.

Similarly, if the attack came from the Internet, you can trace the IP address back to the Internet Service Provider (ISP). ISPs issue IP addresses to users and the ISP logs identify exactly who was issued an IP address at any given time. For example, when David Kernell hacked into Sarah Palin's Yahoo! email account in 2008, security experts quickly traced the attack back to him based on his IP address.

Chapter 8 presented information on logs. Logs record what happened during an event, when it happened, and what account was used during the event. You may remember that a Security log records logon and logoff events. Similarly, many applications require users to authenticate, and

applications log authentication events. All of these logs can be invaluable in re-creating the details of an event after a security incident, including the identity of the account used in the attack.

Chain of Custody

A key part of incident response is collecting and protecting evidence. A chain of custody is a process that provides assurances that evidence has been controlled and handled properly after collection. Forensic experts establish a chain of custody when they first collect evidence.

Security professionals use a chain-of-custody form to document this control. The chain-of-custody form provides a record of every person who was in possession of a physical asset collected as evidence. It shows who had custody of the evidence and where it was stored the entire time since collection. Additionally, personnel often tag the evidence as part of a chain-of-custody process. A proper chain-of-custody process ensures that evidence presented in a court of law is the same evidence that security professionals collected.

As an example, imagine that Homer collected a hard drive as part of an investigation. However, instead of establishing a chain of custody, he simply stores the drive on his desk with the intention of analyzing it the next day. Is it possible that someone could modify the contents of the drive overnight? Absolutely. Instead, he should immediately establish a chain of custody and lock the drive in a secure storage location.

If evidence is not controlled, someone can modify, tamper, or corrupt it. Courts will rule the evidence inadmissible if there is a lack of adequate control, or even a lack of documentation showing that personnel maintained adequate control. However, the chain of custody provides proof that personnel handled the evidence properly.

Remember this

A chain of custody provides assurances that evidence has been controlled and handled properly after collection. It documents who handled the evidence and when they handled it.

Capture Video

Chapter 2 introduced video surveillance methods such as closed-circuit television (CCTV) systems. These can be used as a detective control during an investigation. They provide reliable proof of a person’s location and activity. For example, if a person is stealing equipment or data, video may provide proof.

I remember a high school student was working nights at a local grocery store. The store had a delivery of beer in a tractor-trailer that hadn’t been unloaded yet but was kept backed up to the store loading dock overnight. The student stole several cases of beer thinking the crime was undetectable. However, the entire scene was recorded on video, and when he showed up for work the next evening, the store promptly called the police and provided a copy of the video. The video provided reliable proof that simply couldn’t be disputed.

Record Time Offset

In some cases, it’s easy to identify the time of an event such as in Figure 11.1. In the figure, you can easily identify the exact dates and times when someone created, modified, last saved, and last accessed the file. However, in some cases you need to consider a time offset.




Figure 11.1: Windows Explorer showing exact dates and times

For example, Greenwich Mean Time (GMT) identifies the time at the Royal Observatory in Greenwich, London. Other times are based on GMT. I live in the Eastern Standard Time (EST) zone, so you can express the Date Accessed time as 5:10 p.m. EST. However, GMT uses an offset, so the same time is also 9:10 p.m. GMT. One benefit of using GMT is that it doesn’t change for daylight saving time, so it stays constant.

Many video recorders use time offsets to identify times on tape recordings rather than the actual time. For example, a recording may use a displayed counter to identify the time that has passed since the recording started. Imagine that the counter advances 1,000 ticks or counts per hour. If the counter indicates an event occurred at an offset time of 1,500 and the recording started at midnight, then the time of the event was 1:30 a.m.

When analyzing timestamps of any evidence, it’s important to understand that these times are often based on an offset. If you can’t identify the offset, you may not be able to identify the actual time.

Screenshots

Screenshots are simply pictures of what you can see displayed on a computer screen. If you want to capture exactly what a user was doing, or specific displays, a screenshot is the perfect solution.

For example, Figure 11.1, shown previously, is a screenshot of Windows Explorer. You can save screenshots as graphics files and embed these

graphics into documents. Many operating systems include the ability to capture the screen and save it to the Clipboard. For example, you can capture the screen of almost any system by pressing the PrtScn key found on most keyboards. Many applications such as Snagit allow you to capture screenshots from specific windows or applications, any region of the screen, and even scrolling windows such as a long web page.

Witnesses

Another element of an investigation is interviewing witnesses. Witnesses provide firsthand reports of what happened and when it happened. However, witnesses won't necessarily come forward with relevant information unless someone asks them. Often witnesses don't recognize what information is valuable.

For example, imagine an attacker who tailgated behind an employee without showing credentials. The employee may notice, but not give it much thought, especially if tailgating is common in the organization. If the attack resulted in loss of equipment or data, an investigator may get a good description of the attacker just by interviewing witnesses.

Track Man-Hours and Expense

Investigations can take an extraordinary amount of time, and for any business, time is money. When budget time rolls around, the departments that can accurately identify how much time and money they spent are more likely to get their requested budget approved.

Additionally, quantitative risk assessments base decisions using specific monetary amounts, such as cost and asset values. If an incident required involvement by security professionals on an incident response team, the man-hours and expenses incurred by the incident response team need to be included in the assessment. Including this data improves the accuracy of the cost values used in the quantitative risk assessment.

Big Data Analysis

Big Data is a relatively new term and refers to databases that are so large that tools don't exist to extract meaningful information from them. It's worth stressing that Big Data is really big. It's so big that analysts are currently trying to figure out what to name the next iteration. Currently, yottabyte is the largest name and it refers to 1,000 zetabytes. For the record, the order is gigabyte, terabyte, petabyte, exabyte, zettabyte, and then yottabyte. The next name might be hellabyte, using northern California slang of "hella" meaning "a lot." Seriously. This is one of the names proposed to the International System of Units.

Because tools don't exist to mine this data, analysts must develop tools for specific needs. As an example, SAS Institute developed tools for Macy's Inc. to adjust online pricing for 73 million products based on demand and inventory.

Similarly, forensic specialists often must develop tools to analyze Big Data during forensic investigations. Specialists who understand Big Data and can develop tools to analyze it will become more and more valuable as time goes forward. However, they must be able to develop tools that are reliable and provide consistent results. Similarly, the tools must generate the same results each time they are used against the same data, and be able to preserve the evidence. Clearly, analyzing Big Data brings significant new challenges. For example, it's not a simple matter to create an image of a database as big as a hellabyte.

Raising Security Awareness

Many organizations create a security education and awareness plan to identify methods of raising the security awareness of employees. The primary goal is to minimize the risk posed by users and help to reinforce user compliance with security policies.

Training is especially useful if technical controls are not available to enforce a security policy. For example, if employees are sharing cipher codes for restricted areas, a security control cannot stop them. However, by training the employees of the risks, they are more likely to comply with the security policies.

For example, many users are unaware of the risks associated with USB flash drives. They know that USB flash drives are very convenient and restricting their use sometimes makes it more difficult to do their job. However, they don't always know that an infected USB drive may infect a system as soon as it's plugged in, and an infected system will infect any other USB drives plugged into the system. With a little bit of training, users understand the risks and are more likely to comply with a restrictive USB flash drive policy.

The success of any security awareness and training plan is directly related to the support from senior management. If senior management supports the plan, middle management and employees will also support it. On the other hand, if senior management does not show support for the plan, it's very likely that personnel within the organization will not support it either.

Security Policy Training and Procedures

Organizations often include sections on training and procedures within a security policy. This reminds personnel of the importance of security training and awareness programs. For example, security training isn't a one-time event. Personnel are trained when they are hired and periodically afterwards. For example, it's common to have annual refresher training. This informs personnel of current and updated threats and helps reinforce the importance of user compliance with existing policies.

Additionally, security awareness programs help to keep personnel aware of security risks. Posters and signs help people remember that security is everyone's responsibility. Proxy servers and unified threat management (UTM) devices have URL filters that block access to prohibited sites, such as gambling sites. They typically display a message with the user's name and account information, mention that access was blocked and logged, and remind users of the acceptable use policy.

Security and IT personnel occasionally send out emails when they learn of an emerging threat, such as a tricky phishing attack, or after an incident caused by someone not following established policies. Depending on the extent of the incident, management might require users to complete additional training.

Remember this

A security training and security awareness program helps reduce risks. Security awareness programs educate users about emerging threats and techniques attackers are currently using.

Role-Based Training

Training is often targeted to users based on their roles within the organization. For example, consider the following roles within an organization:

- **Executive personnel.** Executives need high-level briefings related to the risks that the organization faces, along with information on the organization's overall information security awareness program. Additionally, executives should be trained on whaling attacks because attackers target executives with malicious phishing emails.
- **Incident response team.** An incident response team needs detailed training on how to respond to incidents. Even within the team, personnel might require different training. For example, security personnel responsible for forensic investigations need specialized training.
- **Administrators.** Network and server administrators need to understand the hardware and software that they manage, so that they can deploy, manage, and maintain it as securely as possible.
- **End users.** End users need to have an understanding of common threats, such as malware and phishing attacks. They also need to understand the risk posed by clicking an unknown link and how drive-by downloads can infect their system.

Training can include a wide variety of topics depending on the organization. Some of the topics include:

- Security policy contents
- Keeping cipher codes private
- Acceptable use and user responsibilities
- Protection of Personally Identifiable Information
- Importance of data labeling, handling, and disposal
- Information classifications used by the organization
- Compliance with relevant laws, best practices, and standards
- Threat awareness, including current malware and phishing attacks
- User habits that represent risks such as with passwords and tailgating
- Use of social networking sites and peer-to-peer applications and how they result in data leakage

The following section on data policies covers some of these topics, while other topics are covered in previous chapters.

Remember this

A primary goal of security awareness and training is to reinforce user compliance with security policies and help reduce risks posed by users. The success of any security awareness and training plan is dependent on the support of senior management. Because security issues change over time, it's common to provide periodic refresher training.

Training and Compliance Issues

There are many situations where training is required to maintain compliance with existing laws, best practices, and standards. As an example, many laws exist covering PII. Although these laws have many similarities, there can be minor differences in different localities. It's important for personnel handling any PII to understand the laws that apply.

Best practices often prevent a wide range of incidents as long as users understand and follow them. This book has covered many best practices, including developing and following a security policy, ensuring users do not share accounts, using strong passwords, following the principle of least privilege, and much more. Unless personnel know about them, and understand them, they might not be implementing them.

Additionally, many organizations have to abide by certain standards. For example, organizations handling credit card information need to comply with the Payment Card Industry Data Security Standard (PCI DSS). PCI DSS includes six control objectives and 12 specific requirements that combined help ensure an organization implements a series of best practices that will help prevent fraud.

Administrators might understand how to implement many of these without any additional training. However, some of the requirements might require additional training to maintain compliance. PCI DSS isn't foolproof, but it has helped reduce many of the risks associated with credit card fraud.

Using Metrics to Validate Compliance

Metrics measure various activities and in some cases management uses them to measure the impact of training. For example, it's possible for an organization to provide training to personnel and then follow up by gathering metrics. These metrics can validate compliance of personnel following established security policies and measure the overall security posture.

Imagine an organization has been having an average of 10 security incidents a month due to malware and phishing attacks. They might choose to provide training to personnel. Imagine that after personnel attended the training, the organization continued to have an average of 10 security incidents. This indicates the training did not have an impact.

On the other hand, imagine the number of incidents dropped from an average of 10 a month to 1 a month. This indicates that the training was very effective. One of the benefits of these metrics is that they justify the costs associated with training.

Remember this

Metrics can prove the success of a training or security awareness program by comparing incidents before the training with incidents after the training program.

Chapter 11 Exam Topic Review

When preparing for the exam, make sure you understand these key concepts covered in this chapter.

Exploring Security Policies

- Written security policies are management controls that identify an overall security plan for an organization and help to reduce overall risk. Other security controls enforce security policies.
- An acceptable use policy defines proper system usage for users. It often specifically mentions unacceptable usage such as visiting certain web sites, and typically includes statements informing users that the organization monitors user activities. Users are required to read and sign an acceptable use policy when hired, and in conjunction with refresher training.
- Mandatory vacation policies require employees to take time away from their job. These policies help to reduce fraud and discover malicious activities by employees.
- A separation of duties policy separates individual tasks of an overall function between different entities or different people, and helps deter fraud. For example, a single person shouldn't be able to approve bills and pay them, or print checks and then sign them.
- Job rotation policies require employees to change roles on a regular basis. Employees might swap roles temporarily, such as for three to four weeks, or permanently. These policies help to prevent employees from continuing with fraudulent activities, and detect fraud if it occurs.
- Clean desk policies require users to organize their desks and surrounding areas to reduce the risk of possible data theft and password compromise.
- Account policies often require administrators to have two accounts to prevent privilege escalation and other attacks. Account disablement policies ensure that inactive accounts are disabled.
- Change management policies define the process for making changes, and provide the accounting structure or method to document the changes. Change management helps reduce unintended outages from changes.
- Third-party agreements typically include a non-disclosure agreement requiring all parties to recognize who owns the data and prohibiting unauthorized sharing of data.
- A service level agreement (SLA) is an agreement between a company and a vendor that stipulates performance expectations, such as minimum uptime and maximum downtime levels. A memorandum of understanding (MOU) is a looser agreement than an SLA.
- An interconnection security agreement (ISA) specifies technical and security requirements for connections and ensures data confidentiality while data is in transit. An ISA is often used with an MOU.
- Information classification practices help protect sensitive data by ensuring users understand the value of data. Data labeling ensures that users know what data they are handling and processing.
- Degaussing a disk magnetically erases all the data. Physically destroying a drive is the most secure method of ensuring unauthorized personnel cannot access proprietary information.
- Sanitization procedures ensure data is removed from decommissioned systems. Specialized applications erase disk drives by writing a series of 1s and 0s multiple times on the drive. Cluster tip wiping erases file remnants in reclaimed space.
- Storage and retention policies identify how long data is retained. They can limit a company's exposure to legal proceedings and reduce the amount of labor required to respond to court orders.
- Personally Identifiable Information (PII) is used to personally identify an individual. Examples include the full name, birth date, address, and medical information of a person.
- PII requires special handling and policies for data retention. Many laws mandate the protection of PII, and require informing individuals when an attack results in the compromise of PII.
- A privacy policy identifies what data is collected from users on a web site. Many laws require a privacy policy.

Responding to Incidents

- An incident response policy defines an incident and incident response procedures. Organizations review and update incidents periodically and after reviewing lessons learned after actual incidents.
- Warning banners remind users of rules regarding access when the users log on.
- The first step in incident response is preparation. It includes creating and maintaining an incident response policy and includes prevention steps such as implementing security controls to prevent malware infections.
- Before taking action, personnel verify an event is an actual incident. Next, they attempt to contain or isolate the problem. Disconnecting a computer from a network will isolate it.
- First responders are the first IT or security personnel on the scene of an incident. They often have access to toolkits along with contact information of other security personnel.

- An incident response policy typically includes a list of personnel to notify after an incident. A data breach typically requires notifying outside entities, especially if the data breach compromises customer data.
- Recovery or reconstitution restores a system to its original state. Depending on the scope of the incident, administrators might completely rebuild the system, including applying all updates and patches.
- A review of lessons learned helps an organization prevent a recurrence of an incident.
- The order of volatility for data from most volatile to least volatile is cache memory, regular RAM, swap or paging file, hard drive data, logs stored on remote systems, and archived media.
- Forensic experts capture an image of the data before analysis to preserve the original and maintain its usability as evidence.
- Hard drive imaging creates a forensic copy and prevents the forensic capture and analysis from modifying the original evidence. A forensic image is a bit-by-bit copy of the data and does not modify the data during the capture.
- Hashing provides integrity for images, including images of both memory and disk drives. Taking a hash before and after capturing a disk image verifies that the capturing process did not modify data. Hashes can reveal evidence tampering or, at the very least, that evidence has lost integrity.
- A chain of custody provides assurances that personnel controlled and handled evidence properly after collecting it. It may start with a tag attached to the physical item, followed by a chain-of-custody form that documents everyone who handled it and when they handled it.

Raising Security Awareness

- Security awareness and training programs reinforce user compliance with security policies and help reduce risks posed by users.
- Information security awareness programs help educate users about emerging threats such as techniques attackers are currently using, acceptable use policies, and policies related to social networking sites.
- Role-based training ensures that personnel receive the training they need. For example, executives need training on whaling attacks.
- Social media sites allow people to share comments with a wide group of people.
- Improper use of social networking sites can result in inadvertent information disclosure. Attackers gather information from these sites to launch attacks against users, such as cognitive password attacks to change users' passwords. Training reduces these risks.
- Banner ad malware (also known as malvertisements) look like ads but include malicious code. Organizations sometimes block access to some web sites to block banner ad malware.
- Data breaches on social media sites can expose user passwords. If users do not have different passwords, or use the same credentials to access other web applications, a data breach on a social media site can impact much more than just that site.
- P2P software is a source of data leakage. Organizations often block P2P software to prevent data leakage and to prevent P2P traffic from consuming network bandwidth.
- Metrics can validate the success of a training program.

Chapter 11 Practice Questions

1. A security manager needs to identify a policy that will reduce the risk of personnel within an organization colluding to embezzle company funds.

Which of the following is the BEST choice?

- A. AUP
- B. Training
- C. Mandatory vacations
- D. Time-of-day restrictions

2. A security auditor discovered that several employees in the Accounting department can print and sign checks. In her final report, she recommended restricting the number of people who can print checks and the number of people who can sign them. She also recommended that no one should be authorized to print and sign checks. What policy is she recommending?

- A. Discretionary access control
- B. Rule-based access control
- C. Separation of duties
- D. Job rotation

3. Your organization includes a software development division within the IT department. One developer writes and maintains applications for the Sales and Marketing departments. A second developer writes and maintains applications for the Payroll department. Once a year, they have to switch roles for at least a month. What is the purpose of this practice?

- A. To enforce a separation of duties policy
- B. To enforce a mandatory vacation policy
- C. To enforce a job rotation policy
- D. To enforce an acceptable use policy

4. A security manager is reviewing security policies related to data loss. Which of the following is the security administrator MOST likely to be reviewing?

- A. Clean desk policy
- B. Separation of duties
- C. Job rotation
- D. Change management

5. Get Certified Get Ahead (GCGA) has outsourced some application development to your organization. Unfortunately, developers at your organization are having problems getting an application module to work and they want to send the module with accompanying data to a third-party vendor for help in resolving the problem. Which of the following should developers consider before doing so?

- A. Ensure that data in transit is encrypted.
- B. Review NDAs.
- C. Identify the classification of the data.
- D. Verify the third party has an NDA in place.

6. Two companies have decided to work together on a project and implemented an MOU. Which of the following represents the GREATEST security risk in this situation?

- A. An MOU doesn't define responsibilities.
- B. An MOU includes monetary penalties if one party doesn't meet its responsibilities.
- C. An MOU can impose strict requirements for connections.
- D. An MOU doesn't have strict guidelines to protect sensitive data.

7. Your organization is considering storage of sensitive data in a cloud provider. Your organization wants to ensure the data is encrypted while at rest and while in transit. What type of interoperability agreement can your organization use to ensure the data is encrypted while in transit?

- A. SLA
- B. BPA

- C. MOU
- D. ISA

8. A user recently worked with classified data on an unclassified system. You need to sanitize all the reclaimed space on this system's hard drives while keeping the system operational. Which of the following methods will BEST meet this goal?
- A. Use a cluster tip wiping tool.
 - B. Use a file shredding tool.
 - C. Degauss the disk.
 - D. Physically destroy the disk.
9. A network administrator needs to update the operating system on switches used within the network. Assuming the organization is following standard best practices, what should the administrator do first?
- A. Submit a request using the baseline configuration process.
 - B. Submit a request using the incident management process.
 - C. Submit a request using the change management process.
 - D. Submit a request using the application patch management process.
10. Security personnel recently released an online training module advising employees not to share personal information on any social media web sites that they visit. What is this advice MOST likely trying to prevent?
- A. Spending time on non-work-related sites
 - B. Phishing attack
 - C. Cognitive password attacks
 - D. Rainbow table attack
11. Your organization blocks access to social media web sites. The primary purpose is to prevent data leakage, such as the accidental disclosure of proprietary information. What is an additional security benefit of this policy?
- A. Improves employee productivity
 - B. Enables cognitive password attacks
 - C. Prevents P2P file sharing
 - D. Protects against banner ad malware
12. Your organization hosts a web-based server that remote administrators access via Telnet. Management wants to increase their rights to prosecute unauthorized personnel who access this server. Which of the following is the BEST choice?
- A. Enable SSH instead of Telnet.
 - B. Enable banner ads.
 - C. Enable FTP logging.
 - D. Add a warning banner.
13. An incident response team is following typical incident response procedures. Which of the following phases is the BEST choice for analyzing an incident with a goal of identifying steps to prevent a reoccurrence of the incident?
- A. Preparation
 - B. Identification
 - C. Mitigation
 - D. Lessons learned
14. After a recent incident, a forensic analyst was given several hard drives to analyze. What should the analyst do first?
- A. Take screenshots and capture system images.
 - B. Take hashes and screenshots.
 - C. Take hashes and capture system images.
 - D. Perform antivirus scans and create chain-of-custody documents.
15. A forensic expert is preparing to analyze a hard drive. Which of the following should the expert do FIRST?

- A. Capture an image.
- B. Identify the order of volatility.
- C. Create a chain-of-custody document.
- D. Take a screenshot.

16. A security analyst tagged a computer stating when he took possession of it. What is the BEST explanation for this?

- A. To calculate time offset
- B. To ensure the system is decommissioned
- C. To begin a chain of custody
- D. To implement separation of duties

17. You are helping your organization create a security policy for incident response. Of the following choices, what is the BEST choice to include when an incident requires confiscation of a physical asset?

- A. Ensure hashes are taken first.
- B. Ensure witnesses sign an AUP.
- C. Maintain the order of volatility.
- D. Keep a record of everyone who took possession of the physical asset.

18. An administrator recently learned of an attack on a Virginia-based web server from IP address 72.52.206.134 at 11:35:33 GMT. However, after investigating the logs, he is unable to see any traffic from that IP address at that time. Which of the following is the MOST likely reason why the administrator was unable to identify the attack?

- A. He did not account for time offsets.
- B. He did not capture an image.
- C. The IP address has expired.
- D. The logs were erased when the system was rebooted.

19. Personnel in an organization are sharing their access codes to cipher locks with unauthorized personnel. As a result, unauthorized personnel are accessing restricted areas of the building. What is the BEST response to reduce this risk?

- A. Implement a management control.
- B. Implement a technical control.
- C. Implement an AUP.
- D. Provide security training to personnel.

20. Your organization has spent a significant amount of money on training employees on security awareness. Your organization wants to validate the success of this training. Which of the following is the BEST choice?


- A. Implement role-based training.
- B. Use metrics.
- C. Use security policies.
- D. Verify PII.

Performance-Based Question

The “Implementing Basic Forensic Procedures” section of this chapter provides a great example of how you can be tested using a drag-and-drop or matching type of performance-based question. As long as you know the content, these questions typically aren’t any more difficult than a standard multiple-choice question. Here’s an example of a performance-based question.

Instructions: A web server has recently been attacked and a first responder has disconnected it from the network to isolate it. A forensic analyst is preparing to analyze the server but needs to capture data in a specific order to ensure it is preserved. Identify the correct order of the items listed in Figure 11.2.

Server



1

2

3

4

5

Hard Drive

CPU Cache

Remote Logs

RAM

Swap

Figure 11.2: Identifying correct order

Chapter 11 Practice Question Answers

1. **C.** Mandatory vacations help to reduce the possibility of fraud and embezzlement. An acceptable use policy informs users of company policies and even though users sign them, they don't deter someone considering theft by embezzling funds. Training can help reduce incidents by ensuring personnel are aware of appropriate policies. Time-of-day restrictions prevent users from logging on during restricted times.
2. **C.** This recommendation is enforcing a separation of duties principle, which prevents any single person from performing multiple job functions that might allow the person to commit fraud. Discretionary access control specifies that every object has an owner, but doesn't separate duties. Devices such as routers use a rule-based access control model, but it doesn't separate duties. Job rotation policies rotate employees into different jobs, but they don't necessarily separate job functions.
3. **C.** This practice enforces a job rotation policy where employees rotate into different jobs, and is designed to reduce potential incidents. A separation of duties policy prevents any single person from performing multiple job functions to help prevent fraud, but it doesn't force users to switch roles. A mandatory vacation policy requires employees to take time away from their job. An acceptable use policy informs users of their responsibilities when using an organization's equipment.
4. **A.** A clean desk policy requires users to organize their areas to reduce the risk of possible data theft and password compromise. A separation of duties policy separates individual tasks of an overall function between different people. Job rotation policies require employees to change roles on a regular basis. Change management helps reduce intended outages from changes.
5. **B.** Developers should review the non-disclosure agreements (NDAs) and verify that sharing data with a third party doesn't violate any existing NDAs. Encrypting data in transit protects its confidentiality while in transit, but it won't protect it from a third party accessing it after receiving it. The classification of the data isn't as relevant as the NDA in this situation. An NDA between the third party and your organization isn't relevant, if the NDA between you and the hiring organization states you cannot share the data.
6. **D.** A memorandum of understanding (MOU) represents an agreement and it doesn't have strict guidelines to protect sensitive data. An MOU does define responsibilities between the parties. A service level agreement (SLA) might include monetary penalties, but an MOU does not. An interconnection security agreement (ISA) includes strict requirements for connections and is often used with an MOU.
7. **D.** An interconnection security agreement (ISA) specifies technical and security requirements for secure connections and can ensure data is encrypted while in transit. None of the other agreements address the connection. A service level agreement (SLA) stipulates performance expectations of a vendor. A business partners agreement (BPA) is a written agreement for business partners. A memorandum of understanding (MOU) expresses an understanding between two parties to work together.
8. **A.** A cluster tip wiping tool sanitizes reclaimed space on hard drives. The cluster tip is the extra space in the last cluster of a file, which can hold remnants of data. A file shredding tool successfully erases a file, but does not affect clusters in reclaimed space. Degaussing the disk magnetically erases it, and physically destroying the disk is the most secure method protecting its confidentiality, but both of these methods take the system out of operation.
9. **C.** The network administrator should submit a change using the change management process, which is the same process that is typically used for changes to any devices or systems. A baseline configuration identifies the starting configuration. Incident management addresses security incidents. A regular patch management process typically includes following change management, but application patch management does not apply to devices.
10. **C.** A cognitive password attack utilizes information that a person would know, such as the name of their first pet or their favorite color. If this information is available on Facebook or another social media site, attackers can use it to change the user's password. This advice has nothing to do with employees visiting the sites, only with what they post. Although attackers may use this information in a phishing attack, they can also launch phishing attacks without this information. A rainbow table attack is a password attack, but it uses a database of precalculated hashes.
11. **D.** The primary benefit is protection against banner ad malware, also known as malvertisements. Although the policy might result in improved employee productivity, this is not a security benefit. You want to prevent cognitive password attacks, not enable them. Although organizations typically try to prevent peer-to-peer (P2P) file sharing, this is done by blocking access to P2P sites, not social media sites.
12. **D.** A warning banner displayed when personnel log on could inform them that unauthorized access is restricted and is the best choice of those given. Although Secure Shell (SSH) is a more secure alternative than Telnet, it doesn't impact the ability of prosecuting personnel. Banner ads are used on web sites, not within a Telnet session. File Transfer Protocol (FTP) logging wouldn't log Telnet sessions.
13. **D.** You should analyze an incident during the lessons learned stage of incident response with the goal of identifying steps to prevent reoccurrence. Preparation is a planning step done before an incident, with the goal of preventing incidents and identifying methods to respond to incidents. Identification is the first step after hearing about a potential incident to verify it is an incident. Mitigation steps attempt to reduce the effects of the incident.
14. **C.** Forensic analysts capture images and take hashes before beginning analysis, and they only analyze the image copies, not the original drive. Screenshots are taken when a computer is running. An antivirus scan might modify the drive and chain-of-custody documents are created when evidence is collected.
15. **A.** Before analyzing a hard drive, a forensic expert should capture an image of the hard drive and then analyze the image. This protects it from

- accidental modifications and preserves it as usable evidence. The order of volatility identifies what data is most volatile (such as cache) and what is least volatile (such as hard drives). A chain-of-custody document should be created when evidence is first collected. A screenshot is taken when a system is operational.
16. **C.** A chain of custody identifies who controlled evidence after it was confiscated. It can start with a tag when a person collects the evidence. Security analysts later create a chain-of-custody log to detail who controlled the evidence at different times. Time offset is related to different time zones or times recorded on a video recorder. A security analyst would confiscate a computer to analyze it, not decommission it. Separation of duties is related to people, not computers.
17. **D.** It's important to keep a chain of custody for any confiscated physical items and the chain of custody is a record of everyone who took possession of the asset after it was first confiscated. Hashes should be taken before capturing an image, but they are not required before confiscating equipment. Users, not witnesses, sign an acceptable use policy (AUP). Security personnel should be aware of the order of volatility, but there isn't any way to maintain the order.
18. **A.** The most likely reason is that he did not account for the time offset. The attack occurred at 11:35:33 Greenwich Mean Time (GMT) and the web server is in the Eastern Standard Time (EST) zone in Virginia, which is five hours different from GMT. There is no need to capture an image to view logs. IP addresses on the Internet do not expire. Logs are written to a hard drive or a central location; they are not erased when a system is rebooted.
19. **D.** The best response of those listed is to provide training to personnel on the importance of keeping access codes private. Management controls include policies and assessments, but they won't necessarily focus on sharing access codes. Technical controls won't do any good if personnel are bypassing them, which is the case in this scenario. If an acceptable use policy (AUP) isn't implemented, it would be a good idea to implement one. However, it addresses usage of systems, and not necessarily cipher access codes.
20. **B.** Metrics are measurements and you can use them to validate the success of a security awareness program. Role-based training is targeted training, but it does not validate the success of training. Training would typically teach employees about a security policy, but the policy doesn't provide measurements. Personally Identifiable Information (PII) might be part of the training, but PII cannot validate training.

Performance-Based Question Answer

Figure 11.3 shows the solution to the drag-and-drop question and the following text describes the reasoning behind this order:




Figure 11.3: Identifying correct order solution

- Cache memory is the most volatile of all items listed and it should be collected first. This includes central processing unit (CPU) cache or any other type of cache used in the system. Cache typically includes recently used data and information used by applications. It is more volatile than regular RAM because a system has significantly less cache memory than regular RAM, so a system will overwrite cache quicker than regular RAM.
- Random access memory (RAM) is slightly less volatile than cache memory. It can include information used by the system and network processes. It will be lost if the system is powered down (as will the cache memory).
- A swap (or paging file) is an extension of RAM, but it is stored on the hard drive. The swap file is rebuilt each time the system is rebooted, so it is more volatile than regular data stored on a hard drive.
- Data on the hard drive is semipermanent. It remains on the hard drive even after the system is powered down and rebooted.
- Remote logs (or logs stored on remote systems) are less volatile than data stored on the target system. For this reason, many servers send log data to a remote system for centralized collection. Even if the server is completely destroyed, the centralized logs remain unmodified.

CompTIA Security+ Practice Exam

Use this practice exam as an additional study aid before taking the live exam. An answer key with explanation is available at the end of the practice exam.

1. Lisa hid several plaintext documents within an image file. Which security goal is she pursuing?
 - A. Encryption
 - B. Integrity
 - C. Steganography
 - D. Confidentiality
2. You are the security administrator in your organization. You want to ensure that a file maintains integrity. Which of the following choices is the BEST choice to meet your goal?
 - A. Steganography
 - B. Encryption
 - C. Hash
 - D. AES
3. An e-commerce web site does not currently have an account recovery process for customers who have forgotten their passwords. Which of the following choices are the BEST items to include if web site designers add this process? (Select TWO.)
 - A. Create a web-based form that verifies customer identities using another method.
 - B. Set a temporary password that expires upon first use.
 - C. Implement biometric authentication.
 - D. Email the password to the user.
4. Your organization is planning to implement stronger authentication for remote access users. An updated security policy mandates the use of token-based authentication with a password that changes every 30 seconds. Which of the following choices BEST meets this requirement?
 - A. CHAP
 - B. Smart card
 - C. HOTP
 - D. TOTP
5. Your organization issues laptops to mobile users. Administrators configured these laptops with full disk encryption, which requires users to enter a password when they first turn on the computer. After the operating system loads, users are required to log on with a username and password. Which of the following choices BEST describes this?
 - A. Single-factor authentication
 - B. Dual-factor authentication
 - C. Multifactor authentication
 - D. SAML
6. Users at your organization currently use a combination of smart cards and passwords, but an updated security policy requires multifactor security using three different factors. Which of the following can you add to meet the new requirement?
 - A. Four-digit PIN
 - B. Hardware tokens
 - C. Fingerprint readers
 - D. USB tokens
7. A network includes a ticket-granting ticket server used for authentication. What authentication service does this network use?
 - A. TACACS+
 - B. SAML
 - C. LDAP

D. Kerberos

8. You are modifying a configuration file used to authenticate Unix accounts against an external server. The file includes phrases such as DC=Server1 and DC=Com. Which authentication service is the external server using?

- A. Diameter
- B. RADIUS
- C. LDAP
- D. SAML

9. Which of the following choices is an AAA protocol that uses shared secrets as a method of security?

- A. Kerberos
- B. SAML
- C. RADIUS
- D. MD5

10. Your organization wants to reduce the amount of money it is losing due to thefts. Which of the following is the BEST example of an equipment theft deterrent?

- A. Remote wiping
- B. Cable locks
- C. Strong passwords
- D. Disk encryption

11. A manager recently observed an unauthorized person in a secure area, which is protected with a cipher lock door access system. After investigation, he discovered that an authorized employee gave this person the cipher lock code. Which of the following is the BEST response to this issue at the minimum cost?

- A. Implement a physical security control.
- B. Install tailgates
- C. Provide security awareness training.
- D. Place a guard at the entrance.

12. Management recently rewrote the organization's security policy to strengthen passwords created by users. It now states that passwords should support special characters. Which of the following choices is the BEST setting to help the organization achieve this goal?

- A. History
- B. Maximum age
- C. Minimum length
- D. Complexity

13. You have discovered that some users have been using the same passwords for months, even though the password policy requires users to change their password every 30 days. You want to ensure that users cannot reuse the same password. Which settings should you configure? (Select TWO.)

- A. Maximum password age
- B. Password length
- C. Password history
- D. Password complexity
- E. Minimum password age

14. A company recently hired you as a security administrator. You notice that some former accounts used by temporary employees are currently enabled. Which of the following choices is the BEST response?

- A. Disable all the temporary accounts.
- B. Disable the temporary accounts you've noticed are enabled.
- C. Craft a script to identify inactive accounts based on the last time they logged on.
- D. Set account expiration dates for all accounts when creating them.

15. An organization supports remote access, allowing users to work from home. However, management wants to ensure that personnel cannot log on to work systems from home during weekends and holidays. Which of the following BEST supports this goal?

- A. Least privilege
- B. Need to know
- C. Time-of-day restrictions
- D. Mandatory access control

16. You configure access control for users in your organization. Some departments have a high employee turnover, so you want to simplify account administration. Which of the following is the BEST choice?

- A. User-assigned privileges
- B. Group-based privileges
- C. Domain-assigned privileges
- D. Network-assigned privileges

17. You are configuring a file server used to share files and folders among employees within your organization. However, employees should not be able to access all folders on this server. Which of the following choices is the BEST method to manage security for these folders?

- A. Assign permissions to each user as needed.
- B. Wait for users to request permission and then assign the appropriate permissions.
- C. Delegate authority to assign these permissions.
- D. Use security groups with appropriate permissions.

18. The Retirement Castle uses groups for ease of administration and management. They recently hired Jasper as their new accountant. Jasper needs access to all the files and folders used by the Accounting department. What should the administrator do to give Jasper appropriate access?

- A. Create an account for Jasper and add the account to the Accounting group.
- B. Give Jasper the password for the Guest account.
- C. Create an account for Jasper and use rule-based access control for accounting.
- D. Create an account for Jasper and add the account to the Administrators group.

19. Your organization recently updated its security policy and indicated that Telnet should not be used within the network. Which of the following should be used instead of Telnet?

- A. SCP
- B. SFTP
- C. SSL
- D. SSH

20. One of your web servers was recently attacked and you have been tasked with reviewing firewall logs to see if you can determine how an attacker accessed the system remotely. You identified the following port numbers in log entries: 21, 22, 25, 53, 80, 110, 443, and 3389. Which of the following protocols did the attacker MOST likely use?

- A. Telnet
- B. HTTPS
- C. DNS
- D. RDP

21. Which of the following provides the largest address space?

- A. IPv4
- B. IPv5
- C. IPv6
- D. IPv7

22. While analyzing a firewall log, you notice traffic going out of your network on UDP port 53. What does this indicate?

- A. Connection with a botnet
- B. DNS traffic
- C. SMTP traffic
- D. SFTP traffic

23. A team of users in your organization needs a dedicated subnet. For security reasons, other users should not be able to connect to this subnet. Which of the following choices is the BEST solution?

- A. Restrict traffic based on port numbers.
- B. Restrict traffic based on physical addresses.
- C. Implement DNS on the network.
- D. Enable SNMP.

24. An organization recently updated its security policy. A new requirement dictates a need to increase protection from rogue devices plugging into physical ports. Which of the following choices provides the BEST protection?

- A. Disable unused ports
- B. Implement 802.1x
- C. Enable MAC limiting
- D. Enable MAC filtering

25. What would administrators typically place at the end of an ACL of a firewall?

- A. Allow all all
- B. Timestamp
- C. Password
- D. Implicit deny

26. Your organization wants to protect its web server from cross-site scripting attacks. Which of the following choices provides the BEST protection?

- A. WAF
- B. Network-based firewall
- C. Host-based firewall
- D. IDS

27. Management recently learned that several employees are using the company network to visit gambling and gaming web sites. They want to implement a security control to prevent this in the future. Which of the following choices would meet this need?

- A. WAF
- B. UTM
- C. DMZ
- D. NIDS

28. Which of the following protocols operates on Layer 7 of the OSI model?

- A. IPv6
- B. TCP
- C. ARP
- D. SCP

29. Which of the following BEST describes a false negative?

- A. An IDS falsely indicates a buffer overflow attack occurred.
- B. Antivirus software reports that a valid application is malware.
- C. A locked door opens after a power failure.
- D. An IDS does not detect a buffer overflow attack.

30. Company management suspects an employee is stealing critical project information and selling it to a competitor. They'd like to identify who is doing this, without compromising any live data. What is the BEST option to meet this goal?

- A. Install antivirus software on all user systems.
- B. Implement an IPS.
- C. Implement an IDS.
- D. Add fabricated project data on a honeypot.

31. Attackers frequently attack your organization, and administrators want to learn more about zero-day attacks on the network. What can they use?

- A. Anomaly-based HIDS
- B. Signature-based HIDS
- C. Honeypot
- D. Signature-based NIDS

32. Security personnel recently noticed a successful exploit against an application used by many employees at their company. They notified the company that sold them the software and asked for a patch. However, they discovered that a patch wasn’t available. What BEST describes this scenario?

- A. Zero-day
- B. Buffer overflow
- C. LSO
- D. SQL injection

33. What type of encryption is used with WPA2 CCMP?

- A. AES
- B. TKIP
- C. RC4
- D. SSL

34. Administrators in your organization are planning to implement a wireless network. Management has mandated that they use a RADIUS server and implement a secure wireless authentication method. Which of the following should they use?

- A. LEAP
- B. WPA-PSK
- C. WPA2-PSK
- D. AES

35. Which of the following wireless security mechanisms is subject to a spoofing attack?

- A. WEP
- B. IV
- C. WPA2 Enterprise
- D. MAC address filtering

36. Which of the following is the BEST description of why disabling SSID broadcast is not an effective security measure against attackers?

- A. The network name is contained in wireless packets in plaintext.
- B. The passphrase is contained in wireless packets in plaintext.
- C. The SSID is included in MAC filters.
- D. The SSID is not used with WPA2.

37. You are reviewing logs from a wireless survey within your organization’s network due to a suspected attack and you notice the following entries:

MAC	SSID	Encryption	Power
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	47
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	62
56:CD:34:EF:12:AB	GetCertifiedGetAhead	WPA2	20
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	57
12:AB:34:CD:56:EF	GetCertifiedGetAhead	WPA2	49

Of the following choices, what is the MOST likely explanation of these entries?

- A. An evil twin is in place.
- B. Power of the AP needs to be adjusted.
- C. A rogue AP is in place.
- D. The AP is being pharmed.

38. Mobile users in your network report that they frequently lose connectivity with the wireless network on some days, but on other days they don't have any problems. Which of the following types of attacks could cause this?

- A. IV
- B. Wireless jamming
- C. Replay
- D. WPA cracking

39. Management within your organization wants some users to be able to access internal network resources from remote locations. Which of the following is the BEST choice to meet this need?

- A. WAF
- B. VPN
- C. IDS
- D. IPS

40. You suspect that an executable file on a web server is malicious and includes a zero-day exploit. Which of the following steps can you take to verify your suspicion?

- A. Perform a code review.
- B. Perform an architecture review.
- C. Perform a design review.
- D. Perform an operating system baseline comparison.

41. Lisa has scanned all the user computers in the organization as part of a security audit. She is creating an inventory of these systems, including a list of applications running on each computer and the application versions. What is she MOST likely trying to identify?

- A. System architecture
- B. Application baseline
- C. Code vulnerabilities
- D. Attack surface

42. An updated security policy identifies authorized applications for company-issued mobile devices. Which of the following would prevent users from installing other applications on these devices?

- A. Geo-tagging
- B. Authentication
- C. ACLs
- D. Whitelisting

43. A company is implementing a feature that allows multiple servers to operate on a single physical server. What is this?

- A. Virtualization
- B. IaaS
- C. Cloud computing
- D. DLP

44. A software vendor recently developed a patch for one of its applications. Before releasing the patch to customers, the vendor needs to test it in different environments. Which of the following solutions provides the BEST method to test the patch in different environments?

- A. Baseline image

- B. BYOD
- C. Virtualized sandbox
- D. Change management

45. Your company has recently standardized servers using imaging technologies. However, a recent security audit verified that some servers were immune to known OS vulnerabilities, whereas other systems were not immune to the same vulnerabilities. Which of the following would reduce these vulnerabilities?

- A. Patch management
- B. Sandboxing
- C. Snapshots
- D. Baselines

46. Someone stole an executive's smartphone, and the phone includes sensitive data. What should you do to prevent the thief from reading the data?

- A. Password-protect the phone.
- B. Encrypt the data on the phone.
- C. Use remote wipe.
- D. Track the location of the phone.

47. Your organization has issued mobile devices to several key personnel. These devices store sensitive information. What can administrators implement to prevent data loss from these devices if they are stolen?

- A. Inventory control
- B. GPS tracking
- C. Full device encryption
- D. Geo-tagging

48. Homer wants to ensure that other people cannot view data on his mobile device if he leaves it unattended. What should he implement?

- A. Encryption
- B. Cable lock
- C. Screen lock
- D. Remote wiping

49. Management wants to implement a system that will provide automatic notification when personnel remove devices from the building. Which of the following security controls will meet this requirement?

- A. Video monitoring
- B. RFID
- C. Geo-tagging
- D. Account lockout

50. Your organization was recently attacked, resulting in a data breach, and attackers captured customer data. Management wants to take steps to better protect customer data. Which of the following will BEST support this goal?

- A. Succession planning and data recovery procedures
- B. Fault tolerance and redundancy
- C. Stronger access controls and encryption
- D. Hashing and digital signatures

51. A business owner is preparing to decommission a server that has processed sensitive data. He plans to remove the hard drives and send them to a company that destroys them. However, he wants to be certain that personnel at that company cannot access data on the drives. Which of the following is the BEST option to meet this goal?

- A. Encrypt the drives using full disk encryption.
- B. Capture an image of the drives.

- C. Identify data retention policies.
- D. Use file-level encryption to protect the data.

52. Your organization is considering the purchase of new computers. A security professional stresses that these devices should include TPMs. What benefit does a TPM provide? (Choose all that apply.)

- A. It uses hardware encryption, which is quicker than software encryption.
- B. It uses software encryption, which is quicker than hardware encryption.
- C. It includes an HSM file system.
- D. It stores RSA keys.

53. What functions does an HSM include?

- A. Reduces the risk of employees emailing confidential information outside the organization
- B. Provides webmail to clients
- C. Provides full drive encryption
- D. Generates and stores keys used with servers

54. Homer installed code designed to enable his account automatically, three days after someone disables it. What did Homer create?

- A. Backdoor
- B. Rootkit
- C. Armored virus
- D. Ransomware

55. Your local library is planning to purchase new computers that patrons can use for Internet research. Which of the following are the BEST choices to protect these computers? (Choose TWO.)

- A. Mantrap
- B. Anti-malware software
- C. Cable locks
- D. Pop-up blockers
- E. Disk encryption

56. Your organization has been receiving a significant amount of spam with links to malicious web sites. You want to stop the spam. Of the following choices, what provides the BEST solution?

- A. Add the domain to a block list
- B. Use a URL filter
- C. Use a MAC filter
- D. Add antivirus software

57. Attackers have launched an attack using multiple systems against a single target. What type of attack is this?

- A. DoS
- B. DDoS
- C. SYN flood
- D. Buffer overflow

58. Security administrators are reviewing security controls and their usefulness. Which of the following attacks will account lockout controls prevent? (Choose TWO.)

- A. DNS poisoning
- B. Replay
- C. Brute force
- D. Buffer overflow
- E. Dictionary

59. A web developer wants to reduce the chances of an attacker successfully launching XSRF attacks against a web site application. Which of the following provides the BEST protection?
- A. Client-side input validation
 - B. Web proxy
 - C. Antivirus software
 - D. Server-side input validation
60. A web developer is adding input validation techniques to a web site application. Which of the following should the developer implement during this process?
- A. Perform the validation on the server side.
 - B. Perform the validation on the client side.
 - C. Prevent boundary checks.
 - D. Encrypt data with TLS.
61. An attacker is attempting to write more data into a web application's memory than it can handle. What type of attack is this?
- A. XSRF
 - B. LDAP injection
 - C. Fuzzing
 - D. Buffer overflow
62. During a penetration test, a tester injected extra input into an application causing the application to crash. What does this describe?
- A. SQL injection
 - B. Fuzzing
 - C. Transitive access
 - D. XSRF
63. A security expert is attempting to identify the number of failures a web server has in a year. Which of the following is the expert MOST likely identifying?
- A. SLE
 - B. MTTR
 - C. ALE
 - D. MTTF
64. You are trying to add additional security controls for a database server that includes customer records and need to justify the cost of \$1,000 for these controls. The database includes 2,500 records. Estimates indicate a cost of \$300 for each record if an attacker successfully gains access to them. Research indicates that there is a 10 percent possibility of a data breach in the next year. What is the ALE?
- A. \$300
 - B. \$37,500
 - C. \$75,000
 - D. \$750,000
65. A penetration tester is tasked with gaining information on one of your internal servers and he enters the following command: **telnet server1 80**. What is the purpose of this command?
- A. Identify if server1 is running a service using port 80 and is reachable.
 - B. Launch an attack on server1 sending 80 separate packets in a short period of time.
 - C. Use Telnet to remotely administer server1.
 - D. Use Telnet to start an RDP session.
66. A recent vulnerability assessment identified several issues related to an organization's security posture. Which of the following issues is MOST likely to affect the organization on a day-to-day basis?

- A. Natural disasters
- B. Lack of antivirus software
- C. Lack of protection for data at rest
- D. Lack of protection for data in transit

67. Which of the following tools would a security administrator use to identify misconfigured systems within a network?

- A. Pentest
- B. Virus scan
- C. Load test
- D. Vulnerability scan

68. A security expert is running tests to identify the security posture of a network. However, these tests are not exploiting any weaknesses. Which of the following types of test is the security expert performing?

- A. Penetration test
- B. Virus scan
- C. Port scan
- D. Vulnerability scan

69. Which of the following tools is the LEAST invasive and can verify if security controls are in place?

- A. Pentest
- B. Protocol analyzer
- C. Vulnerability scan
- D. Host enumeration

70. Your organization develops web application software, which it sells to other companies for commercial use. To ensure the software is secure, your organization uses a peer assessment to help identify potential security issues related to the software. Which of the following is the BEST term for this process?

- A. Code review
- B. Change management
- C. Routine audit
- D. Rights and permissions review

71. Your organization plans to deploy new systems within the network within the next six months. What should your organization implement to ensure these systems are developed properly?

- A. Code review
- B. Design review
- C. Baseline review
- D. Attack surface review

72. You need to periodically check the configuration of a server and identify any changes. What are you performing?

- A. Code review
- B. Design review
- C. Attack surface review
- D. Baseline review

73. Your organization hired an external security expert to test a web application. The security expert is not given any access to the application interfaces, code, or data. What type of test will the security expert perform?

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

74. A security administrator needs to inspect protocol headers of traffic sent across the network. What tool is the BEST choice for this task?
- A. Web security gateway
 - B. Protocol analyzer
 - C. Honeypot
 - D. Vulnerability assessment
75. You are troubleshooting issues between two servers on your network and need to analyze the network traffic. Of the following choices, what is the BEST tool to capture and analyze this traffic?
- A. Switch
 - B. Protocol analyzer
 - C. Firewall
 - D. NIDS
76. Which of the following is the lowest cost solution for fault tolerance?
- A. Load balancing
 - B. Clustering
 - C. RAID
 - D. Cold site
77. You need to modify the network infrastructure to increase availability of web-based applications for Internet clients. Which of the following choices provides the BEST solution?
- A. Load balancing
 - B. Proxy server
 - C. UTM
 - D. Content inspection
78. A security analyst is creating a document that includes the expected monetary loss from a major outage. She is calculating the potential lost sales, fines, and impact on the organization's customers. Which of the following documents is she MOST likely creating?
- A. BCP
 - B. BIA
 - C. DRP
 - D. RPO
79. Your organization is updating its business continuity documents. You're asked to review the communications plans for possible updates. Which of the following should you ensure is included in the communications plan?
- A. A list of systems to recover in hierarchical order
 - B. Incident response procedures
 - C. List of critical systems and components
 - D. Methods used to respond to media requests, including templates
80. What type of encryption does the RADIUS protocol use?
- A. Symmetric
 - B. Asymmetric
 - C. MD5
 - D. SHA
81. Your organization is planning to implement videoconferencing, but it wants to protect the confidentiality of the streaming video. Which of the following would BEST meet this need?
- A. PBKDF2
 - B. DES
 - C. MD5

D. RC4

82. An organization is implementing a PKI and plans on using public and private keys. Which of the following can be used to create strong key pairs?

- A. MD5
- B. RSA
- C. AES
- D. HMAC

83. Your organization is investigating possible methods of sharing encryption keys over a public network. Which of the following is the BEST choice?

- A. CRL
- B. PBKDF2
- C. Hashing
- D. ECDHE

84. A user wants to hide confidential data within a .jpg file. Which of the following is the BEST choice to meet this need?

- A. ECC
- B. Steganography
- C. CRL
- D. File-level encryption

85. You need to ensure data sent over an IP-based network remains confidential. Which of the following provides the BEST solution?

- A. Stream ciphers
- B. Block ciphers
- C. Transport encryption
- D. Hashing

86. Personnel within your company are assisting an external auditor perform a security audit. They frequently send documents to the auditor via email and some of these documents contain confidential information. Management wants to implement a solution to reduce the possibility of unintentionally exposing this data. Which of the following is the BEST choice?

- A. Hash all outbound email containing confidential information.
- B. Use digital signatures on all outbound email containing confidential information.
- C. Encrypt all outbound email containing confidential information.
- D. Implement DLP to scan all outbound email.

87. Which two protocols provide strong security for the Internet with the use of certificates? (Choose TWO.)

- A. SSH
- B. SSL
- C. SCP
- D. TLS
- E. SFTP

88. Lenny and Carl work in an organization that includes a PKI. Carl needs to send a digitally signed file to Lenny. What does Carl use in this process?

- A. Carl's public key
- B. Carl's private key
- C. Lenny's public key
- D. Lenny's private key

89. Bart recently sent out confidential data via email to potential competitors. Management suspects he did so accidentally, but Bart denied sending the data. Management wants to implement a method that would prevent Bart from denying accountability in the future. What are they trying to enforce?

- A. Confidentiality
- B. Encryption
- C. Access control

D. Non-repudiation

90. An organization is planning to implement an internal PKI for smart cards. Which of the following should the organization do FIRST?

- A. Install a CA.
- B. Generate key pairs.
- C. Generate a certificate.
- D. Identify a recovery agent.

91. Which of the following is a valid reason to use a wildcard certificate?

- A. Reduce the administrative burden of managing certificates.
- B. Support multiple private keys.
- C. Support multiple public keys.
- D. Increase the lifetime of the certificate.

92. Homer works as a contractor at a company on a one-year renewing contract. After renewing his contract, the company issues him a new smart card. However, he is now having problems digitally signing email or opening encrypted email. What is the MOST likely solution?

- A. Copy the original certificate to the new smart card.
- B. Copy his original private key to the new smart card.
- C. Copy his original public key to the new smart card.
- D. Publish the certificate in his new smart card.

93. You need to request a certificate for a web server. Which of the following would you MOST likely use?

- A. CA
- B. CRL
- C. CSR
- D. OCSP

94. An organization is implementing a data policy and wants to designate a recovery agent. Which of the following indicates what a recovery agent can do?

- A. A recovery agent can retrieve a user's public key.
- B. A recovery agent can decrypt data if users lose their private key.
- C. A recovery agent can encrypt data if users lose their private key.
- D. A recovery agent can restore a system from backups.

95. An organizational policy specifies that duties of application developers and administrators must be separated. What is the MOST likely result of implementing this policy?

- A. One group develops program code and the other group deploys the code.
- B. One group develops program code and the other group modifies the code.
- C. One group deploys program code and the other group administers databases.
- D. One group develops databases and the other group modifies databases.

96. Application developers in your organization currently update applications on live production servers when needed. However, they do not follow any predefined procedures before applying the updates. What should the organization implement to prevent any risk associated with this process?

- A. Risk assessment
- B. Tabletop exercises
- C. Change management
- D. Incident management

97. Which of the following is a type of media that allows the mass distribution of personal comments to specific groups of people?

- A. P2P
- B. Social media

- C. Media devices
- D. News media

98. Your organization wants to prevent damage from malware. Which stage of the common incident response procedures is the BEST stage to address this?
- A. Preparation
 - B. Identification
 - C. Mitigation
 - D. Lessons learned
99. You are reviewing incident response procedures related to the order of volatility. Which of the following is the LEAST volatile?
- A. Hard disk drive
 - B. Memory
 - C. RAID-6 cache
 - D. CPU cache
100. Security personnel confiscated a user's workstation after a security incident. Administrators removed the hard drive for forensic analysis, but left it unattended for several hours before capturing an image. What could prevent the company from taking the employee to court over this incident?
- A. Witnesses were not identified.
 - B. A chain of custody was not maintained.
 - C. An order of volatility was not maintained.
 - D. A hard drive analysis was not complete.

Security+ Practice Exam Answers

When checking your answers, take the time to read the explanation. Understanding the explanations will help ensure you're prepared for the live exam. The explanations also show the chapter or chapters where you can get more detailed information on the topic.

1. **D.** Hiding files in another file is one way to achieve the security goal of confidentiality. In this scenario, Lisa is using steganography as the method by hiding files within a file. Encryption is the best way to achieve confidentiality, but simply hiding files within a file doesn't encrypt the data. Hashing methods and digital signatures provide integrity. See Chapters 1 and 10.
2. **C.** A hash provides integrity for files, emails, and other types of data. Steganography provides confidentiality by hiding data within other data and encryption provides confidentiality by ciphering the data. Advanced Encryption Standard (AES) is an encryption protocol. See Chapter 1.
3. **A, B.** A web-based form using an identity-proofing method, such as requiring users to enter the name of their first pet, can verify their identity. Setting a password that expires upon first use ensures that the user changes the password. Biometric authentication is not reasonable for an online e-commerce web site. Emailing the password is a possibility, but not without configuring the password to expire upon first use. See Chapter 1.
4. **D.** A Time-based One-Time Password (TOTP) creates passwords that expire after 30 seconds. An HMAC-based One Time Password (HOTP) creates passwords that do not expire. Challenge Handshake Authentication Protocol uses a nonce (a number used once), but a nonce does not expire after 30 seconds. See Chapter 1.
5. **A.** Both passwords are in the something you know factor of authentication, so this process is single-factor authentication. Dual-factor authentication requires the use of two different authentication factors. Multifactor authentication requires two or more factors of authentication. Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML) used for single sign-on (SSO), but this is unrelated to this question. See Chapter 1.
6. **C.** Fingerprint readers would add biometrics from the something you are factor of authentication as a third factor of authentication. The current system includes methods in the something you have factor (smart cards) and in the something you know factor (passwords), so any solution requires a method that isn't using one of these two factors. A PIN is in the something you know factor. Hardware tokens and USB tokens are in the something you have factor. See Chapter 1.
7. **D.** Kerberos uses a ticket-granting ticket server, which creates tickets for authentication. Terminal Access Controller Access-Control System Plus (TACACS+) is an authentication service created by Cisco. Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML) used for single sign-on (SSO) solutions. Lightweight Directory Access Protocol (LDAP) is an X.500-based authentication service that can be secured with Transport Layer Security (TLS). See Chapter 1.
8. **C.** Lightweight Directory Access Protocol (LDAP) uses X.500-based phrases to identify components such as the domain component (DC). Diameter is an alternative to Remote Authentication Dial-In User Service (RADIUS), but neither of these use X.500-based phrases. Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML) used for web-based single sign-on (SSO) solutions. See Chapter 1.
9. **C.** Remote Authentication Dial-In User Service (RADIUS) is an authentication, authorization, and accounting (AAA) protocol that uses shared secrets (or passwords) for security. Kerberos uses tickets. SAML provides SSO for web-based applications, but it is not an AAA protocol. MD5 is a hashing protocol, not an AAA protocol. See Chapter 1.
10. **B.** Cable locks are effective equipment theft deterrents for laptops and other systems. Remote wiping can erase data on stolen systems, but it doesn't deter thefts. Strong passwords help prevent someone from accessing a stolen device, but it doesn't deter thefts. Disk encryption can protect the data after a device is stolen, but it doesn't deter theft. See Chapter 2.
11. **C.** Security awareness training is often the best response to violations of security policies. If individuals do not abide by the policies after training, management can take disciplinary action. The cipher lock is a physical security control, but it is not effective due to employees bypassing it. Tailgating occurs when one user follows closely behind another user without using credentials and mantraps prevent tailgating, but tailgates are on the back of trucks. Guards can prevent this issue by only allowing authorized personnel in based on facial recognition or identification badges, but at a much higher cost. See Chapter 2 and 11.
12. **D.** The complexity setting is the best answer because it includes using multiple character types, such as special characters, numbers, and uppercase and lowercase letters. The history setting remembers previous passwords and prevents users from reusing them. The maximum age setting forces users to change their password after a set number of days has passed. The minimum length setting forces users to create passwords with a minimum number of characters, such as eight. See Chapter 2.
13. **C, E.** The password history setting records previously used passwords (such as the last 24 passwords) to prevent users from reusing the same passwords. Using password history setting combined with the minimum password age setting prevents users from changing their password repeatedly to get back to their original password. The maximum password age setting ensures users change their passwords regularly, but this is already set to 30 days in the scenario. Password length requires a minimum number of characters in a password. Password complexity requires a mix of uppercase and lowercase letters, numbers, and special characters. See Chapter 2.
14. **C.** Running a last logon script allows you to identify inactive accounts, such as accounts that haven't been logged on to in the last 30 days. It's appropriate to disable unused accounts, but it isn't necessarily appropriate to disable all temporary accounts, because some might still be in use. If you

- disable the accounts you notice, you might disable accounts that some employees are still using, and you might miss some accounts that should be disabled. Setting expiration dates for newly created accounts is a good step, but it doesn't address previously created accounts. See Chapter 2.
15. **C.** Time-of-day restrictions prevent users from logging on during certain times. Least privilege and need to know restrict access to only what the user needs, and these concepts are not associated with time. Mandatory access control uses labels and can restrict access based on need to know, but it is not associated with time. See Chapter 2.
16. **B.** Group-based privileges is a form of role-based access control and it simplifies administration. Instead of assigning permissions to new employees individually, you can just add new employee user accounts into the appropriate groups to grant them the rights and permissions they need for the job. User-assigned privileges require you to manage privileges for each user separately, and it increases the account administration burden. Domain-assigned and network-assigned privileges are not valid administration practices. See Chapter 2.
17. **D.** You can create security groups, place users into these groups, and grant access to the folders by assigning appropriate permissions to the security groups. For example, the security groups might be Sales, Marketing, and HR, and you place users into the appropriate group based on their job. This is an example of using group-based privileges. Waiting for users to ask, and then assigning permissions to users individually has a high administrative overhead. Although delegating authority to assign permissions might work, it doesn't provide the same level of security as centrally managed groups, and without groups, it will still have a high administrative overhead for someone. See Chapter 2.
18. **A.** The administrator should create an account for Jasper and add it to the Accounting group. Because the organization uses groups, it makes sense that they have an Accounting group. The Guest account should be disabled to prevent the use of generic accounts. This scenario describes role-based access control, not rule-based access control. Jasper does not require administrator privileges, so his account should not be added to the Administrators group. See Chapter 2.
19. **D.** Secure Shell (SSH) is a good alternative to Telnet. SSH encrypts transmissions, whereas Telnet transmits data in cleartext. Secure Copy (SCP) and Secure File Transfer Protocol (SFTP) use SSH to encrypt files sent over the network. See Chapter 3.
20. **D.** The attacker most likely used Remote Desktop Protocol (RDP) over port 3389. Telnet can connect to systems remotely, but it uses port 23 and that isn't one of the listed ports. HTTPS uses port 443 for secure HTTP sessions. DNS uses port 53 for name resolution queries and zone transfers. See Chapter 3.
21. **C.** Internet Protocol version 6 provides the largest address space using 128 bits to define an IP address. IPv4 uses 32 bits. IPv5 uses 64 bits but was never adopted. IPv7 has not been defined. See Chapter 3.
22. **B.** Domain Name System (DNS) traffic uses UDP port 53 by default to resolve host names to IP addresses. It is not malicious traffic connecting to a botnet. Simple Mail Transfer Protocol (SMTP) uses port 25. Secure File Transfer Protocol (SFTP) uses port 22. See Chapter 3.
23. **B.** Of the given choices, the best answer is to restrict traffic based on physical addresses. This is also known as media access control (MAC) address filtering and is configured on a switch. Port numbers are related to protocols, so it wouldn't be feasible to restrict traffic for this group based on protocols. Domain Name System (DNS) provides name resolution, but it doesn't restrict traffic. Simple Network Management Protocol version 3 (SNMPv3) monitors and manages network devices. See Chapter 3.
24. **B.** IEEE 802.1x is a port-based authentication protocol and it requires systems to authenticate before they are granted access to the network. If an attacker plugged a rogue device into a physical port, the 802.1x server would block it from accessing the network. Disabling unused ports is a good practice, but it doesn't prevent an attacker from unplugging a system from a used port and plugging the rogue device into the port. While MAC limiting and filtering will provide some protection against rogue devices, an 802.1x server provides much stronger protection. See Chapter 3.
25. **D.** Administrators would place an implicit deny rule at the end of an access control list (ACL) to deny all traffic that hasn't been explicitly allowed. Many firewalls place this rule at the end by default. An allow all rule explicitly allows all traffic and defeats the purpose of a firewall. Timestamps aren't needed in an ACL. ACLs are in cleartext so should not include passwords. See Chapter 3.
26. **A.** A web application firewall (WAF) is an Application layer firewall designed specifically to protect web servers. Although both host-based and network-based firewalls provide protection, they aren't necessarily Application layer firewalls, so they do not provide the same level of protection for a web server as a WAF does. An intrusion detection system (IDS) can help detect attacks, but it isn't as good as the WAF when protecting the web server. See Chapter 3.
27. **B.** A unified threat management (UTM) device typically includes a URL filter and can block access to web sites, just as a proxy server can block access to web sites. A web application firewall (WAF) protects a web server from incoming attacks. A demilitarized zone (DMZ) is a buffered zone between protected and unprotected networks, but it does not include URL filters. A network-based intrusion detection system (NIDS) can detect attacks, but doesn't include outgoing URL filters. See Chapter 3.
28. **D.** Secure Copy (SCP) operates on Layer 7 of the OSI model. IPv6 operates on Layer 3. TCP operates on Layer 4. Address Resolution Protocol (ARP) operates on Layer 2. See Chapter 3.
29. **D.** An intrusion detection system (IDS) should detect a buffer overflow attack and report it, but if it does not, it is a false negative. If the IDS falsely indicates an attack occurred, it is a false positive. If antivirus software indicates a valid application is malware, it is a false positive. A locked door that opens after a power failure is designed to fail-open. See Chapter 4.
30. **D.** Fabricated data on a honeypot could lure the malicious insider and entice him to access it. Antivirus software blocks malware. An intrusion

prevention system (IPS) and an intrusion detection system (IDS) each detect attacks, but won't detect someone accessing data on a server. See Chapter 4.

31. **C.** A honeypot is a server designed to look valuable to an attacker and can help administrators learn about zero-day exploits, or previously unknown attacks. A host-based intrusion detection system (HIDS) protects host systems, but isn't helpful against network attacks. Signature-based tools would not have a signature for zero-day attack because the attack method is unknown by definition. See Chapter 4.

32. **A.** This scenario describes a zero-day exploit on the software application. A zero-day exploit is one that is unknown to the vendor, or the vendor knows about, but hasn't yet released a patch or update to mitigate the threat. The other answers are specific types of attacks, but the scenario isn't specific enough to identify the type of exploit. A buffer overflow attack occurs when an attacker attempts to write more data into an application's memory than it can handle, or to bypass the application's structured exception handling (SEH). Adobe Flash content within web pages uses locally shared objects (LSOs), similar to how regular web pages use cookies, and attackers can modify both cookies and LSOs in different types of attacks. A Structured Query Language (SQL) injection attack attempts to inject SQL code into an application to access a database. See Chapter 4.

33. **A.** Wi-Fi Protected Access II (WPA2) with Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP) uses Advanced Encryption Standard (AES). Temporal Key Integrity Protocol (TKIP) and Secure Sockets Layer (SSL) both use Rivest Cipher 4 (RC4), but not AES. See Chapter 4.

34. **A.** Enterprise mode implements 802.1x as a Remote Authentication Dial-In User Service (RADIUS) server and Lightweight Extensible Authentication Protocol (LEAP) can secure the authentication channel. LEAP is a Cisco proprietary protocol, but other EAP variations can also be used, such as Protected EAP (PEAP), EAP-Transport Layer Security (EAP-TLS), and EAP Tunneled TLS (EAP-TTLS). Wi-Fi Protected Access (WPA) and WPA2 using a preshared key (PSK) do not use RADIUS. Many security protocols use Advanced Encryption Standard (AES), but AES by itself does not use RADIUS. See Chapter 4.

35. **D.** Media access control (MAC) address filtering is vulnerable to spoofing attacks because attackers can easily change MAC addresses on network interface cards (NICs). Wired Equivalent Privacy (WEP) can be cracked using an initialization vector (IV) attack, but not by spoofing. WPA2 Enterprise requires users to enter credentials, so it isn't susceptible to a spoofing attack. See Chapter 4.

36. **A.** The service set identifier (SSID) is the network name and it is included in certain wireless packets in plaintext. Disabling SSID broadcast hides the wireless network from casual users, but not attackers. Passphrases are not sent across the network in plaintext and are unrelated to the SSID. Media access control (MAC) address filters do not include the SSID. Wi-Fi Protected Access II (WPA2) does use the SSID. See Chapter 4.

37. **A.** The logs indicate an evil twin is in place. An evil twin is a rogue wireless access point with the same service set identifier (SSID) as a live wireless access point. The SSID is GetCertifiedGetAhead and most of the entries are from an access point (AP) with a media access control (MAC) address of 12:AB:34:CD:56:EF. However one entry shows a MAC of 56:CD:34:EF:12:AB, indicating an evil twin with the same name as the legitimate AP. Power can be adjusted if necessary to reduce the visibility of the AP, but there isn't any indication this is needed. The power of the evil twin is lower, indicating it is in a different location farther away. A rogue AP is an unauthorized AP and although the evil twin is unauthorized, it is more correct to identify this as an evil twin because that is more specific. Generically, a rogue AP has a different SSID. A pharming attack redirects a web site's traffic to another web site, but this isn't indicated in this question at all. See Chapter 4.

38. **B.** A wireless jamming attack is a type of denial-of-service (DoS) attack that can cause wireless devices to lose their association with access points and disconnect them from the network. None of the other attacks are DoS attacks. An initialization vector (IV) is a specific type of attack on Wired Equivalent Privacy (WEP) to crack the key. A replay attack captures traffic with the goal of replaying it later to impersonate one of the parties in the original transmission. Wi-Fi Protected Access (WPA) cracking attacks attempt to discover the passphrase. See Chapter 4.

39. **B.** A virtual private network (VPN) provides access to a private network over a public network such as the Internet via remote locations and is the best choice. A web application firewall (WAF) provides protection for a web application or a web server. Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) protect networks, but do not control remote access. See Chapter 4.

40. **D.** An operating system baseline comparison is the best choice of the available answers. It can verify if the file is in the baseline, or was added after the server was deployed. A code review is possible if you have access to the original code, but this isn't easily possible with an executable file. Code reviews look at the code before it is released and architecture reviews look at architecture designs, but neither of these identifies malicious files after a web server has been deployed. See Chapter 5.

41. **B.** Administrators create a list of applications installed on systems as part of an application baseline (also called a host software baseline). An architecture review typically looks at the network architecture, not individual systems. A code review looks for vulnerabilities within code, but applications are compiled so the code is not easily available for review. The attack surface looks at much more than just applications and includes protocols and services. See Chapter 5.

42. **D.** Whitelisting identifies authorized software and prevents users from installing or running any other software. Geo-tagging adds location information to media such as photographs, but the scenario only refers to applications. Authentication allows users to prove their identity, such as with a username and password, but isn't relevant in this question. Access control lists (ACLs) are used with routers, firewalls, and files, but do not restrict installation of applications. See Chapter 5.

43. **A.** Virtualization allows multiple virtual servers to exist on a single physical server. Infrastructure as a Service (IAAS) is a cloud computing option where the vendor provides access to a computer, but customers manage it. Cloud computing refers to accessing computing resources via a different

location than your local computer. Data loss prevention (DLP) techniques examine and inspect data looking for unauthorized data transmissions. See Chapter 5.

44. **C.** A virtualized sandbox provides a simple method of testing patches and would be used with snapshots so that the virtual machine (VM) can easily be reverted to the original state. A baseline image is a starting point of a single environment. Bring your own device (BYOD) refers to allowing employee-owned mobile devices in a network, and is not related to this question. Change management practices ensure changes are not applied until they are approved and documented. See Chapter 5.

45. **A.** Patch management procedures ensure operating systems (OS) are kept up to date with current patches. Patches ensure systems are immune to known vulnerabilities, but none of the other answers protects systems from these known vulnerabilities. Sandboxing isolates systems for testing. Snapshots record the state of a virtual machine at a moment in time. Baselines identify the starting point for systems. See Chapter 5.

46. **C.** Remote wipe capabilities can send a remote wipe signal to the phone to delete all the data on the phone, including any cached data. The phone is lost, so it's too late to password-protect or encrypt the data now if these steps weren't completed previously. Although tracking the phone might be useful, it doesn't prevent the thief from reading the data. See Chapter 5.

47. **C.** Full device encryption helps prevent data loss in the event of theft of a mobile device storing sensitive information. Other security controls (not listed as answers in this question) that help prevent loss of data in this situation are a screen lock, account lockout, and remote wipe capabilities. Inventory control methods help ensure devices aren't lost or stolen. Global positioning system (GPS) tracking helps locate the device. Geo-tagging includes geographical information with pictures posted to social media sites. See Chapter 5.

48. **C.** A screen lock locks a device until the proper passcode is entered and prevents access to mobile devices when they are left unattended. Encryption protects data, especially if the device is lost or stolen. A cable lock is used with laptops to prevent them from being stolen. Remote wiping can erase data on a lost or stolen device. See Chapter 5.

49. **B.** Radio-frequency identification (RFID) provides automated inventory control and can detect movement of devices. Video monitoring might detect removal of devices, but it does not include automatic notification. Geo-tagging provides geographic location for pictures posted to social media sites. Account lockout controls lock accounts when the incorrect password is entered too many times. See Chapter 5.

50. **C.** Strong access controls and encryption are two primary methods of protecting the confidentiality of any data, including customer data. Succession planning and data recovery procedures are part of business continuity. Fault tolerance and redundancy increase the availability of data. Hashing and digital signatures provide integrity. See Chapter 5.

51. **A.** Full disk encryption is the best option of the available answers. Another option (not listed) is to use disk wiping procedures to erase the data. Capturing an image of the drives won't stop someone from accessing data on the original drives. Retention policies identify how long to keep data, but do not apply here. Depending on how much data is on the drives, file-level encryption can be very tedious and won't necessarily encrypt all of the sensitive data. See Chapter 5.

52. **A, D.** A Trusted Platform Module (TPM) is a hardware chip that stores RSA encryption keys and uses hardware encryption, which is quicker than software encryption. A TPM does not use software encryption. An HSM is a removable hardware device that uses hardware encryption, but it does not have a file system and TPM does not provide HSM as a benefit. See Chapter 5.

53. **D.** A hardware security module (HSM) is a removable device that can generate and store RSA keys used with servers for data encryption. A data loss prevention (DLP) device is a device that can reduce the risk of employees emailing confidential information outside the organization. Software as a Service (SaaS) provides software or applications, such as webmail, via the cloud. A Trusted Platform Module (TPM) provides full drive encryption and is included in many laptops. See Chapter 5.

54. **A.** By ensuring that his account is automatically reenabled, Homer has created a backdoor. He is creating this with a logic bomb, but a logic bomb isn't available as a choice in this question. Rootkits include hidden processes, but they do not activate in response to events. An armored virus uses techniques to make it difficult for researchers to reverse engineer it. Ransomware demands payment to release a user's computer or data. See Chapter 6.

55. **B, C.** Anti-malware software and cable locks are the best choices to protect these computers. Anti-malware software protects the systems from viruses and other malware. The cable locks deter theft of the computers. A mantrap prevents tailgating, but this is unrelated to this question. Pop-up blockers are useful, but they are often included with anti-malware software, so anti-malware software is most important. Disk encryption is useful if the computers have confidential information, but it wouldn't be appropriate to put confidential information on a public computer. See Chapters 2 and 6.

56. **A.** You can block emails from a specific domain sending spam by adding the domain to a block list. While the question doesn't indicate that the spam is coming from a single domain, this is still the best answer of the given choices. A URL filter blocks outgoing traffic and can be used to block the links to the malicious web sites in this scenario, but it doesn't stop the email. Switches use MAC filters to restrict access within a network. Antivirus software does not block spam. See Chapter 6.

57. **B.** A distributed denial-of-service (DDoS) attack includes attacks from multiple systems with the goal of depleting the target's resources. A DoS attack comes from a single system and a SYN flood is an example of a DoS attack. A buffer overflow is a type of DoS attack that attempts to write data into an application's memory. See Chapter 7.

58. **C, E.** Brute force and dictionary attacks attempt to guess passwords, but an account lockout control locks an account after the wrong password is guessed too many times. The other attacks are not password attacks, so they aren't mitigated using account lockout controls. Domain name system (DNS)

poisoning attempts to redirect web browsers to malicious URLs. Replay attacks attempt to capture packets to impersonate one of the parties in an online session. Buffer overflow attacks attempt to overwhelm online applications with unexpected code or data. See Chapter 7.

59. **D.** Validating and filtering input using server-side input validation can restrict the use of special characters needed in cross-site request forgery (XSRF) attacks. Both server-side and client-side input validation is useful, but client-side input validation can be bypassed, so it should not be used alone. A web proxy can filter URLs, but it cannot validate data. Additionally, web proxies can be used to bypass client-side input validation techniques. Antivirus software cannot detect XSRF attacks. See Chapter 7.

60. **A.** Input validation should be performed on the server side. Client-side validation can be combined with server-side validation, but it can be bypassed so it should not be used alone. Boundary or limit checks are an important part of input validation. Input validation does not require encryption of data with Transport Layer Security (TLS) or any other encryption protocol. See Chapter 7.

61. **D.** One type of buffer overflow attack attempts to write more data into an application's memory than it can handle. A cross-site request forgery (XSRF) attack attempts to launch attacks with HTML code. Lightweight Directory Application Protocol (LDAP) injection attacks attempt to query directory service databases such as Microsoft Active Directory. Fuzzing inputs random data into an application during testing. See Chapter 7.

62. **B.** Fuzzing or fuzz testing sends extra input to an application to test it. Ideally, the application can handle the extra input, but it is possible that fuzz testing causes an application to crash. Other answers do not cause the application to crash. A SQL injection attack sends specific SQL code to access or modify data in a database. A cross-site request forgery (XSRF) attack uses HTML or JavaScript code to take actions on behalf of a user. See Chapter 7.

63. **C.** Annualized loss expectancy (ALE) is part of a quantitative risk assessment and is the most likely answer of those given. It is calculated by multiplying the single loss expectancy times the annualized rate of occurrence (ARO). Mean time to recover (MTTR) and mean time to failure (MTTF) do not identify the number of failures in a year. See Chapter 8.

64. **C.** The annual loss expectancy (ALE) is \$75,000. The single loss expectancy (SLE) is \$750,000 (\$300 per record \times 2,500 records). The annual rate of occurrence (ARO) is 10 percent or .10. You calculate the ALE as $SLE \times ARO$ (\$750,000 \times .10). One single record is \$300, but if an attacker can gain access to the database, the attacker can access all 2,500 records. If the ARO was .05, the ALE would be \$37,500. See Chapter 8.

65. **A.** This command sends a query to server1 over port 80 and if the server is running a service on port 80, it will connect. This is a common beginning command for a banner grabbing attempt. It does not send 80 separate packets. If **80** was omitted, Telnet would attempt to connect using its default port of 23 and attempt to create a Telnet session. Remote Desktop Protocol (RDP) uses port 3389 and is not relevant in this scenario. See Chapter 8.

66. **B.** Malware is a constant threat and without antivirus software, systems are sure to become infected in a short period of time. Natural disasters are a risk, but not on a day-to-day basis. Encryption protects data at rest and data in transit, but a lack of encryption isn't likely to affect the organization on a day-to-day basis. See Chapter 8.

67. **D.** A vulnerability scan checks systems for potential vulnerabilities, including vulnerabilities related to misconfiguration. Although a penetration test (pentest) can identify misconfigured systems, it also attempts to exploit vulnerabilities on these systems, so it isn't appropriate if you only want to identify the systems. A virus scan identifies malware and a load test determines if a system can handle a load, but neither of these identifies misconfigured systems. See Chapter 8.

68. **D.** A vulnerability scan identifies the security posture of a network but it does not actually exploit any weaknesses. In contrast, a penetration test attempts to exploit weaknesses. A virus scan searches a system for malware and a port scan identifies open ports, but neither identifies the security posture of an entire network. See Chapter 8.

69. **C.** A vulnerability scan can verify if security controls are in place, and it does not try to exploit these controls using any invasive methods. A pentest (or penetration test) can verify if security controls are in place, but it is invasive and can potentially compromise a system. A protocol analyzer is not invasive, but it cannot determine if security controls are in place. Host enumeration identifies hosts on a network, but does not check for security controls. See Chapter 8.

70. **A.** Peers, such as other developers, perform code reviews going line-by-line through the software code looking for vulnerabilities, such as buffer overflows and race conditions. Change management helps prevent unintended outages from configuration changes. Routine audits review processes and procedures, but not software code. A user rights and permissions review ensures users have appropriate privileges. See Chapter 8.

71. **B.** A design review ensures that systems and software are developed properly. A code review is appropriate if the organization is developing its own software for these new systems, but the scenario doesn't indicate this. A baseline review identifies changes from the initial baseline configuration, but couldn't be done for systems that aren't deployed yet. Identifying the attack surface, including the required protocols and services, would likely be part of the design review, but the design review does much more. See Chapter 8.

72. **D.** A baseline review identifies changes from the original deployed configuration. The original configuration is also known as the baseline. A code review checks internally developed software for vulnerabilities. A design review verifies the design of software or applications to ensure they are developed properly. Determining the attack surface is an assessment technique, but it does not identify changes. See Chapter 8.

73. **D.** A black box tester doesn't have access to any data prior to a test and this includes application interfaces, code, and data. White box testers would be given full access to the application interfaces, code, and data, and gray box testers would be given some access. Black hat refers to a malicious attacker. See Chapter 8.

74. **B.** A protocol analyzer (or sniffer) can capture traffic allowing an administrator to inspect the protocol headers. A web security gateway is a type of

security appliance that protects against multiple threats, but doesn't necessarily capture traffic for inspection. A honeypot contains fake data designed to entice attackers. A vulnerability assessment identifies a system or network's security posture and it might include using a protocol analyzer, but does much more. See Chapter 8.

75. **B.** A protocol analyzer (also called a sniffer) is the best choice to capture and analyze network traffic. Although the traffic probably goes through a switch, the switch doesn't capture the traffic in such a way that you can analyze it. It's unlikely that the traffic is going through a firewall between two internal servers and even if it did, the best you could get is data from the firewall log, but this wouldn't provide the same level of detail as a capture from the sniffer. A network intrusion detection system (NIDS) detects traffic, but it isn't the best tool to capture and analyze it. See Chapter 8.

76. **C.** A redundant array of inexpensive disks (RAID) subsystem is a relatively low-cost solution for fault tolerance for disks. RAID also increases data availability. Load balancing and failover clustering add in additional servers, which is significantly more expensive than RAID. A cold site is a completely separate location, which can be expensive, but a cold site does not provide fault tolerance. See Chapter 9.

77. **A.** Load-balancing solutions increase the availability of web-based solutions by spreading the load among multiple servers. A proxy server is used by internal clients to access Internet resources and does not increase availability of a web server. A unified threat management (UTM) system protects internal resources from attacks, but does not directly increase the availability of web-based applications. Content inspection is one of the features of a UTM, and it protects internal clients but does not directly increase the availability of web-based applications. See Chapter 9.

78. **B.** A business impact analysis (BIA) includes information on potential monetary losses and is the most likely document of those listed that would include this information. A business continuity plan (BCP) includes a BIA, but the BIA is more likely to include this information than the BCP is. A disaster recovery plan (DRP) includes methods used to recover from an outage. The recovery point objective (RPO) refers to the amount of data you can afford to lose but does not include monetary losses. See Chapter 9.

79. **D.** A communications plan will include methods used to respond to media requests, including basic templates. Although not available as a possible answer, it would also include methods used to communicate with response team members, employees, suppliers, and customers. None of the other answers are part of a communications plan. A DRP includes a list of systems to recover in hierarchical order. An incident response plan identifies incident response procedures. A BIA identifies critical systems and components. See Chapter 9.

80. **A.** Remote Authentication Dial-In User Service (RADIUS) uses symmetric encryption. It does not use asymmetric encryption, which uses a public key and a private key. Message Digest 5 (MD5) and Secure Hash Algorithm (SHA) are hashing algorithms.

81. **D.** Rivest Cipher 4 (RC4) is a symmetric encryption stream cipher, and a stream cipher is often the best choice for encrypting data of an unknown size, such as streaming video. Encryption is the best way to ensure the confidentiality of data. Password-Based Key Derivation Function 2 (PBKDF2) is a key stretching technique designed to protect passwords against brute force attempts and is not used for streaming data. Data Encryption Standard (DES) is an older block cipher that is not secure. Message Digest 5 (MD5) is a hashing algorithm used for integrity. See Chapter 10.

82. **B.** Rivest, Shamir, Adleman (RSA) is used to create key pairs. Message Digest 5 (MD5) and Hash-based Message Authentication Code (HMAC) are hashing algorithms. Advanced Encryption Standard (AES) is a symmetric encryption algorithm. See Chapter 10.

83. **D.** Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) allows entities to negotiate encryption keys securely over a public network. Password-Based Key Derivation Function 2 (PBKDF2) is a key stretching technique designed to make password cracking more difficult. A certificate revocation list (CRL) identifies revoked certificates and is unrelated to sharing encryption keys. Hashing methods do not support sharing encryption keys over a public network. See Chapter 10.

84. **B.** Steganography allows users to hide data within the white space of other files, including .jpg files. None of the other choices hides data within another file. Elliptic curve cryptography (ECC) is often used with mobile devices for encryption because it has minimal overhead. A certificate revocation list (CRL) identifies revoked certificates. File-level encryption encrypts a file, such as a master password list, but does not hide data within another file. See Chapter 10.

85. **C.** Transport encryption techniques such as Internet Protocol security (IPsec) provide confidentiality. Both stream ciphers and block ciphers can be used by different transport encryption protocols. Hashing provides integrity, but encryption is needed to provide confidentiality. See Chapters 3, 4, and 10.

86. **C.** The best method of preventing unintentional exposure of confidential information is encryption, so encrypting all outbound emails containing confidential information is the best choice. Hashing the emails doesn't protect the confidentiality of the information. Digital signatures provide proof of who sent an email, but don't protect confidentiality. Data loss prevention (DLP) techniques can detect when employees send out some types of data, but block the transmission and would prevent the auditors from getting the data they need. See Chapter 10.

87. **B, D.** Secure Sockets Layer (SSL) and Transport Layer Security (TLS) secure Internet traffic with the use of certificates. Secure Shell (SSH) encrypts traffic such as Secure Copy (SCP), Secure File Transfer Protocol (SFTP), and Telnet but none of these use certificates. See Chapter 10.

88. **B.** Carl uses his private key to digitally sign the file. Lenny uses Carl's public key to decrypt the digital signature. Lenny's keys are not used in this scenario. See Chapter 10.

89. **D.** Non-repudiation methods such as digital signatures prevent users from denying they took an action. Encryption methods protect confidentiality. Access control methods protect access to data. See Chapters 1 and 10.

90. **A.** A Public Key Infrastructure (PKI) requires a certification authority (CA), so a CA should be installed first. Smart cards require certificates and

would be issued by the CA. After installing the CA, you can generate key pairs to be used with certificates issued by the CA. A recovery agent can be identified, but it isn't required to be done as a first step for a CA. See Chapter 10.

91. **A.** A wildcard certificate reduces the certificate management burden by using an asterisk (*) in place of child domain names. The certificate still has a single public and private key pair. The wildcard doesn't affect the lifetime of the certificate. See Chapter 10.

92. **D.** He should publish the certificate in his new smart card in a global address list within the domain. It is not possible for users to copy a certificate, a public key, or a private key to a smart card. See Chapter 10.

93 **C.** A certificate signing request (CSR) uses a specific format to request a certificate. You submit the CSR to a Certificate Authority (CA), but the request needs to be in the CSR format. A certificate revocation list (CRL) is a list of revoked certificates. The Online Certificate Status Protocol (OCSP) is an alternate method of validating certificates and indicates if a certificate is good, revoked, or unknown. See Chapter 10.

94. **B.** Recovery agents can decrypt data and messages if users lose their private key. Public keys are publicly available, so recovery agents aren't needed to retrieve them. A recovery agent wouldn't encrypt a user's data. Although backups are important, this isn't the role of a recovery agent. See Chapter 10.

95. **A.** This describes a separation of duties policy where the application developers create and modify the code, and the administrators deploy the code to live production systems, but neither group can perform both functions. Developers would typically develop the original code, and modify it when necessary. This scenario does not mention databases. See Chapter 11.

96. **C.** A change management process ensures that changes are approved before being implemented and would prevent risks associated with unintended outages. A risk assessment identifies risks at a given point in time. Tabletop exercises test business continuity and disaster recovery plans. Incident management is only related to security incidents. See Chapter 11.

97. **B.** Social media is a type of media that allows the mass distribution of personal comments to specific groups of people and it is a potential risk to organizations due to possible data leakage. Peer-to-peer (P2P) sites allow users to share data, but it is also a source of data leakage. Media devices such as MP3 players don't support sharing comments among specific groups of users. The news media reports on news stories. See Chapter 11.

98. **A.** The preparation stage is the first phase of common incident response procedures, and attempts to prevent incidents and plan methods to respond to incidents. Incident identification occurs after a potential incident occurs and verifies it is an incident. You attempt to reduce or remove the effects of an incident during the mitigation stage. Lessons learned occurs later and involves analysis to identify steps that will prevent a future occurrence. See Chapter 11.

99. **A.** Data on a hard disk drive is the least volatile of those listed. All other sources are some type of memory, which will be lost if a system is turned off. This includes data in a redundant array of inexpensive disks 6 (RAID-6) cache, normal memory, and the central processing unit's (CPU's) memory. See Chapter 11.

100. **B.** A chain of custody was not maintained because the hard drive was left unattended for several hours before capturing an image. Witnesses were not mentioned, but are not needed if the chain of custody was maintained. The order of volatility does not apply here, but the hard drive is not volatile. Analysis would occur after capturing an image, but there isn't any indication it wasn't done or wasn't complete. See Chapter 11.

Appendix A—Acronym List

This acronym list provides you with a quick reminder of many of the different security-related terms along with a short explanation. Where appropriate, the concepts are explained in greater depth within the book. You can use the index to identify the specific pages where the topics are covered.

802.1x—A port-based authentication protocol. Wireless can use 802.1x. For example, WPA2 Enterprise mode uses an 802.1x server (implemented as a RADIUS server). Enterprise mode requires an 802.1x server. PEAP and EAP-TTLS require a certificate on the 802.1x server. EAP-TLS also uses TLS, but it requires certificates on both the 802.1x server and each of the clients.

3DES—Triple Digital Encryption Standard. A symmetric algorithm used to encrypt data and provide confidentiality. It is a block cipher that encrypts data in 64-bit blocks. It was originally designed as a replacement for DES, and is still used in some applications, such as when hardware doesn't support AES.

AAA—Authentication, Authorization, and Accounting. AAA protocols are used in remote access systems. For example, TACACS+ is an AAA protocol that uses multiple challenges and responses during a session. Authentication verifies a user's identification. Authorization determines if a user should have access. Accounting tracks a user's access with logs.

ACE—Access Control Entry. Identifies a user or group that is granted permission to a resource. ACEs are contained within a DACL in NTFS.

ACK—Acknowledge. A packet in a TCP handshake. In a SYN flood attack, attackers send the SYN packet, but don't complete the handshake after receiving the SYN/ACK packet.

ACL—Access control list. Routers and packet-filtering firewalls perform basic filtering using an ACL to control traffic based on networks, subnets, IP addresses, ports, and some protocols. In NTFS, a list of ACEs makes up the ACL for a resource.

AES—Advanced Encryption Standard. A symmetric algorithm used to encrypt data and provide confidentiality. AES is a block cipher and it encrypts data in 128-bit blocks. It is quick, highly secure, and used in a wide assortment of cryptography schemes. It includes key sizes of 128 bits, 192 bits, or 256 bits.

AES-256—Advanced Encryption Standard 256 bit. AES sometimes includes the number of bits used in the encryption keys and AES-256 uses 256-bit encryption keys. Interestingly, Blowfish is quicker than AES-256.

AH—Authentication Header. IPsec includes both AH and ESP. AH provides authentication and integrity using HMAC. ESP provides confidentiality, integrity, and authentication using HMAC, and AES or 3DES. AH is identified with protocol ID number 51.

ALE—Annual (or annualized) loss expectancy. The ALE identifies the expected annual loss and is used to measure risk with ARO and SLE in a quantitative risk assessment. The calculation is $SLE \times ARO = ALE$.

AP—Access point, short for wireless access point (WAP). APs provide access to a wired network to wireless clients. Many APs support Isolation mode to segment wireless users from other wireless users.

API—Application Programming Interface. A software module or component that identifies inputs and outputs for an application.

APT—Advanced persistent threat. A group that has both the capability and intent to launch sophisticated and targeted attacks.

ARO—Annual (or annualized) rate of occurrence. The ARO identifies how many times a loss is expected to occur in a year and it is used to measure risk with ALE and SLE in a quantitative risk assessment. The calculation is $SLE \times ARO = ALE$.

ARP—Address Resolution Protocol. Resolves IPv4 addresses to MAC addresses. ARP poisoning attacks can redirect traffic through an attacker's system by sending false MAC address updates. NDP is used with IPv6 instead of ARP.

ASCII—American Standard Code for Information Interchange. Code used to display characters.

ASP—Application Service Provider. Provides an application as a service over a network.

AUP—Acceptable use policy. An AUP defines proper system usage. It will often describe the purpose of computer systems and networks, how users can access them, and the responsibilities of users when accessing the systems.

BAC—Business Availability Center. An application that shows availability and performance of applications used or provided by a business.

BCP—Business continuity plan. A plan that helps an organization predict and plan for potential outages of critical services or functions. It includes disaster recovery elements that provide the steps used to return critical functions to operation after an outage. A BIA is a part of a BCP and the BIA drives decisions to create redundancies such as failover clusters or alternate sites.

BIA—Business impact analysis. The BIA identifies systems and components that are essential to the organization's success. It identifies various scenarios that can impact these systems and components, maximum downtime limits, and potential losses from an incident. The BIA helps identify RTOs and RPOs.

BIND—Berkeley Internet Name Domain. BIND is DNS software that runs on Linux and Unix servers. Most Internet-based DNS servers use BIND.

BIOS—Basic Input/Output System. A computer's firmware used to manipulate different settings such as the date and time, boot drive, and access

password. UEFI is the designated replacement for BIOS.

BPA—Business partners agreement. A written agreement that details the relationship between business partners, including their obligations toward the partnership.

BYOD—Bring your own device. A policy allowing employees to connect personally owned devices, such as tablets and smartphones, to a company network. Data security is often a concern with BYOD policies and organizations often use VLANs to isolate mobile devices.

CA—Certificate Authority. An organization that manages, issues, and signs certificates and is part of a PKI. Certificates are an important part of asymmetric encryption. Certificates include public keys along with details on the owner of the certificate and on the CA that issued the certificate. Certificate owners share their public key by sharing a copy of their certificate.

CAC—Common Access Card. A specialized type of smart card used by the U.S. Department of Defense. It includes photo identification and provides confidentiality, integrity, authentication, and non-repudiation for the users. It is similar to a PIV.

CAN—Controller Area Network. A standard that allows microcontrollers and devices to communicate with each other without a host computer.

CAPTCHA—Completely Automated Public Turing Test to Tell Computers and Humans Apart. Technique used to prevent automated tools from interacting with a web site. Users must type in text, often from a slightly distorted image.

CAR—Corrective Action Report. A report used to document actions taken to correct an event, incident, or outage.

CCMP—Counter Mode with Cipher Block Chaining Message Authentication Code Protocol. An encryption protocol based on AES and used with WPA2 for wireless security. It is more secure than TKIP, which was used with the original release of WPA.

CCTV—Closed-circuit television. This is a detective control that provides video surveillance. Video surveillance provides reliable proof of a person's location and activity. It is also a physical security control and it can increase the safety of an organization's assets.

CERT—Computer Emergency Response Team. A group of experts who respond to security incidents. Also known as CIRT, SIRT, or IRT.

CHAP—Challenge Handshake Authentication Protocol. Authentication mechanism where a server challenges a client. More secure than PAP and uses PPP. MS-CHAPv2 is an improvement over CHAP and uses mutual authentication.

CIA—Confidentiality, integrity, and availability. These three form the security triad. Confidentiality helps prevent the unauthorized disclosure of data. Integrity provides assurances that data has not been modified, tampered with, or corrupted. Availability indicates that data and services are available when needed.

CIO—Chief Information Officer. A “C” level executive position in some organizations. A CIO focuses on using methods within the organization to answer relevant questions and solve problems.

CIRT—Computer Incident Response Team. A group of experts who respond to security incidents. Also known as CERT, SIRT, or IRT.

COOP—Continuity of operations planning. Continuity of operations planning (COOP) sites provide an alternate location for operations after a critical outage. A hot site includes personnel, equipment, software, and communication capabilities of the primary site with all the data up to date. A cold site will have power and connectivity needed for COOP activation, but little else. A warm site is a compromise between a hot site and a cold site. Mobile sites do not have dedicated locations, but can provide temporary support during a disaster.

CP—Contingency planning. Plans for contingencies in the event of a disaster to keep an organization operational. BCPs include contingency planning.

CRC—Cyclical Redundancy Check. An error detection code used to detect accidental changes that can affect the integrity of data.

CRL—Certification revocation list. A list of certificates that a CA has revoked. Certificates are commonly revoked if they are compromised, or issued to an employee who has left the organization. The Certificate Authority (CA) that issued the certificate publishes a CRL, and a CRL is public.

CSR—Certificate signing request. A method of requesting a certificate from a CA. It starts by creating an RSA-based private/public key pair and then including the public key in the CSR.

CSR—Control Status Register. A register in a processor used for temporary storage of data.

CSU—Channel Service Unit. A line bridging device used with T1 and similar lines. It typically connects with a DSU as a CSU/DSU.

CTO—Chief Technology Officer. A “C” level executive position in some organizations. CTOs focus on technology and evaluate new technologies.

CVE—Common Vulnerabilities and Exposures (CVE). A dictionary of publicly known security vulnerabilities and exposures.

DAC—Discretionary access control. An access control model where all objects have owners and owners can modify permissions for the objects (files and folders). Microsoft NTFS uses the DAC model. Other access control models are MAC and RBAC.

DACL—Discretionary access control list. List of Access Control Entries (ACEs) in Microsoft NTFS. Each ACE includes a security identifier (SID) and a permission.

DBA—Database administrator. A DBA administers databases on database servers.

dBd—Decibels-dipole. Identifies the gain of an antenna compared with a type of dipole antenna. Higher dBd numbers indicate the antenna can transmit and receive over greater distances.

dBi—Decibels-isotropic. Identifies the gain of an antenna and is commonly used with omnidirectional antennas. It references an isotropic antenna that can theoretically transmit the signal equally in all directions. Higher numbers indicate the antenna can transmit and receive over greater distances.

dBm—Decibels-milliwatt. Identifies the power level of the WAP and refers to the power ratio in decibels referenced to one milliwatt. Higher numbers indicate the WAP transmits the signal over a greater distance.

DDoS—Distributed denial-of-service. An attack on a system launched from multiple sources intended to make a computer’s resources or services unavailable to users. DDoS attacks typically include sustained, abnormally high network traffic. Compare to DoS.

DEP—Data Execution Prevention. A security feature in some operating systems. It helps prevent an application or service from executing code from a nonexecutable memory region.

DES—Digital Encryption Standard. An older symmetric encryption standard used to provide confidentiality. DES is a block cipher and it encrypts data in 64-bit blocks. DES uses 56 bits and is considered cracked. Use AES instead, or 3DES if the hardware doesn’t support AES.

DHCP—Dynamic Host Configuration Protocol. A service used to dynamically assign TCP/IP configuration information to clients. DHCP is often used to assign IP addresses, subnet masks, default gateways, DNS server addresses, and much more.

DHE—Data-Handling Electronics. Term used at NASA indicating electronic systems that handle data.

DHE—Diffie-Hellman Ephemeral. An alternative to traditional Diffie-Hellman. Instead of using static keys that stay the same over a long period, DHE uses ephemeral keys, which change for each new session. Sometimes listed as EDH.

DLL—Dynamic Link Library. A compiled set of code that can be called from other programs.

DLP—Data loss prevention. A network-based DLP system can examine and analyze network traffic. It can detect if confidential company data or any PII data is included in email and reduce the risk of internal users emailing sensitive data outside the organization. End-point DLP systems can prevent users from copying or printing sensitive data.

DMZ—Demilitarized zone. A buffer zone between the Internet and an internal network. It allows access to services while segmenting access to the internal network. Internet clients can access the services hosted on servers in the DMZ, but the DMZ provides a layer of protection for the internal network.

DNAT—Dynamic Network Address Translation. A form of NAT that uses multiple public IP addresses. In contrast, PAT uses a single public IP address. It hides addresses on an internal network.

DNAT—Destination Network Address Translation. A form of NAT that changes the destination IP address for incoming traffic. It is used for port forwarding.

DNS—Domain Name System. Used to resolve host names to IP addresses. DNS zones include records such as A records for IPv4 addresses and AAAA records for IPv6 addresses. DNS uses UDP port 53 for DNS client queries and TCP port 53 for zone transfers. DNS poisoning attacks attempt to modify or corrupt DNS data. Secure zone transfers help prevent these attacks. A pharming attack is a type of DNS poisoning attack that redirects a web site’s traffic to another web site.

DNSSEC—Domain Name System Security Extensions. A suite of specifications used to protect the integrity of DNS records and prevent DNS poisoning attacks.

DoS—Denial-of-service. An attack from a single source that attempts to disrupt the services provided by the attacked system. Compare to DDoS.

DRP—Disaster recovery plan. A document designed to help a company respond to disasters, such as hurricanes, floods, and fires. It includes a hierarchical list of critical systems and often prioritizes services to restore after an outage. Testing validates the plan. The final phase of disaster recovery includes a review to identify any lessons learned and may include an update of the plan.

DSA—Digital Signature Algorithm. A digital signature is an encrypted hash of a message. The sender’s private key encrypts the hash of the message to create the digital signature. The recipient decrypts the hash with the sender’s public key, and, if successful, it provides authentication, non-repudiation, and integrity. Authentication identifies the sender. Integrity verifies the message has not been modified. Non-repudiation is used with online transactions and prevents the sender from later denying he sent the email.

DSL—Digital subscriber line. Improvement over traditional dial-up to access the Internet.

DSU—Data Service Unit. An interface used to connect equipment to a T1 and similar lines. It typically connects with a CSU as a CSU/DSU.

EAP—Extensible Authentication Protocol. An authentication framework that provides general guidance for authentication methods. Variations include EAP-TLS, EAP-TTLS, LEAP, and PEAP.

EAP-TLS—Extensible Authentication Protocol-Transport Layer Security. An extension of EAP sometimes used with 802.1x. This is one of the most secure EAP standards and is widely implemented. The primary difference between PEAP and EAP-TLS is that EAP-TLS requires certificates on the 802.1x server and on each of the wireless clients.

EAP-TTLS—Extensible Authentication Protocol-Tunneled Transport Layer Security. An extension of EAP sometimes used with 802.1x. It allows systems to use some older authentication methods such as PAP within a TLS tunnel. It requires a certificate on the 802.1x server but not on the clients.

ECC—Elliptic curve cryptography. An asymmetric encryption algorithm commonly used with smaller wireless devices. It uses smaller key sizes and requires less processing power than many other encryption methods.

ECDHE—Elliptic Curve Diffie-Hellman Ephemeral. A version of Diffie-Hellman that uses ECC to generate encryption keys. Ephemeral keys are re-created for each session.

EFS—Encrypting File System. A feature within NTFS on Windows systems that supports encrypting individual files or folders for confidentiality.

EMI—Electromagnetic interference. Interference caused by motors, power lines, and fluorescent lights. EMI shielding prevents outside interference sources from corrupting data and prevents data from emanating outside the cable.

ESD—Electrostatic discharge. Release of static electricity. ESD can damage equipment and low humidity causes a higher incidence of electrostatic discharge (ESD). High humidity can cause condensation on the equipment, which causes water damage.

ESN—Electronic Serial Number. Numbers used to uniquely identify mobile devices.

ESP—Encapsulating Security Protocol. IPsec includes both AH and ESP. AH provides authentication and integrity using HMAC. ESP provides confidentiality, integrity, and authentication using HMAC and AES or 3DES. ESP is identified with protocol ID number 50.

FACL—File System Access Control List. An ACL used for file systems. As an example, NTFS uses the DAC model to protect files and folders.

FCoE—Fibre Channel over Ethernet. A lower-cost alternative to traditional SANs. It supports sending Fibre Channel commands over an IP network.

FDE—Full Disk Encryption. Method to encrypt an entire disk. TrueCrypt is an example.

FTP—File Transfer Protocol. Used to upload and download files to an FTP server. FTP uses TCP ports 20 and 21. Secure FTP (SFTP) uses SSH for encryption on TCP port 22. FTP Secure (FTPS) uses SSL or TLS for encryption.

FTPS—File Transfer Protocol Secure. An extension of FTP that uses SSL to encrypt FTP traffic. Some implementations of FTPS use TCP ports 989 and 990.

GPG—GNU Privacy Guard (GPG). Free software based on the OpenPGP standard and used to encrypt and decrypt files. It is similar to PGP but avoids any conflict with existing licensing by using open standards.

GPO—Group Policy Object. Group Policy is used within Microsoft Windows to manage users and computers. It is implemented on a domain controller within a domain. Administrators use it to create password policies, lock down the GUI, configure host-based firewalls, and much more.

GPS—Global Positioning System. GPS tracking can help locate lost mobile devices. Remote wipe, or remote sanitize, erases all data on lost devices. Full disk encryption protects the data on the device if it is lost.

GRE—Generic Routing Encapsulation. A tunneling protocol developed by Cisco Systems.

GUI—Graphical user interface. Users interact with the graphical elements instead of typing in commands from a text interface. Windows is an example of a GUI.

HDD—Hard disk drive. A disk drive that has one or more platters and a spindle. In contrast, USB flash drives and SSD drives use flash memory.

HIDS—Host-based intrusion detection system. An IDS used to monitor an individual server or workstation. It protects local resources on the host such as the operating system files, and in some cases, it can detect malicious activity missed by antivirus software.

HIPS—Host-based intrusion prevention system. An extension of a host-based IDS. Designed to react in real time to catch an attack in action.

HMAC—Hash-based Message Authentication Code. A hashing algorithm used to verify integrity and authenticity of a message with the use of shared secret. When used with TLS and IPsec, HMAC is combined with MD5 and SHA-1 as HMAC-MD5 and HMAC-SHA1, respectively.

HOTP—[HMAC-based One-Time Password](#) (HOTP). An open standard used for creating one-time passwords, similar to those used in tokens or key fobs. It combines a secret key and an incrementing counter, and then uses HMAC to create a hash of the result. HOTP passwords do not expire until they are used.

HSM—Hardware security module. A removable or external device that can generate, store, and manage RSA keys used in asymmetric encryption. High-volume e-commerce sites use HSMs to increase the performance of SSL sessions. High-availability clusters needing encryption services can use clustered HSMs.

HTML—Hypertext Markup Language. Language used to create web pages. HTML documents are displayed by web browsers and delivered over the Internet using HTTP or HTTPS. It uses less-than and greater-than characters (< and >) to create tags. Many sites use input validation to block these tags and prevent cross-site scripting attacks.

HTTP—Hypertext Transfer Protocol. Used for web traffic on the Internet and in intranets. HTTP uses TCP port 80.

HTTPS—Hypertext Transfer Protocol Secure. Encrypts HTTP traffic with SSL or TLS using TCP port 443.

HVAC—Heating, ventilation, and air conditioning. HVAC systems increase availability by regulating airflow within data centers and server rooms. They use hot and cold aisles to regulate the cooling, thermostats to ensure a relatively constant temperature, and humidity controls to reduce the potential for static discharge, and damage from condensation. Higher-tonnage HVAC systems provide more cooling capacity to keep server rooms at operating temperatures, resulting in fewer failures and longer MTBF times. HVAC systems should be integrated with fire alarm systems and either have dampers or the ability to be turned off in the event of a fire.

IaaS—Infrastructure as a Service. A cloud computing technology that allows an organization to rent access to hardware. It provides customers with access to hardware in a self-managed platform. Customers are responsible for keeping an IaaS system up to date. Compare to PaaS and SaaS.

ICMP—Internet Control Message Protocol. Used for diagnostics such as ping. Many DoS attacks use ICMP. It is common to block ICMP at firewalls and routers. If ping fails, but other connectivity to a server succeeds, it indicates that ICMP is blocked.

ID—Identification. For example, a protocol ID identifies a protocol based on a number. AH is identified with protocol ID number 51 and ESP is identified with protocol ID number 50.

IDS—Intrusion detection system. A detective control used to detect attacks after they occur. Monitors a network (NIDS) or host (HIDS) for intrusions and provides ongoing protection against various threats. IDSs include sniffing capabilities. Many IDSs use numbering systems to identify vulnerabilities.

IEEE—Institute of Electrical and Electronics Engineers. IEEE is an international organization with a focus on electrical, electronics, and information

technology topics. IEEE standards are well respected and followed by vendors around the world.

IGMP—Internet Group Management Protocol. Used for multicasting. Computers belonging to a multicasting group have a multicasting IP address in addition to a standard unicast IP address.

IIS—Internet Information Services. A Microsoft Windows web server. IIS comes free with Microsoft Windows Server products. Linux systems use Apache as a web server.

IKE—Internet Key Exchange. Used with IPsec to create a secure channel over UDP port 500 in a VPN tunnel.

IM—Instant messaging. Real-time direct text-based communication between two or more people, often referred to as chat. Spim is a form of spam using IM.

IMAP4—Internet Message Access Protocol v4. Used to store email on servers and allow clients to manage their email on the server. IMAP4 uses TCP port 143.

IP—Internet Protocol. Used for addressing. See IPv4 and IPv6.

IPS—Intrusion prevention system. A preventive control that will stop an attack in progress. It is similar to an active IDS except that it's placed in-line with traffic. An IPS can actively monitor data streams, detect malicious content, and stop attacks in progress. It can be used internally to protect private networks, such as those holding SCADA equipment.

IPsec—Internet Protocol security. Used to encrypt data in transit and can operate in both Tunnel mode and Transport mode. It uses Tunnel mode for VPN traffic. IPsec is built in to IPv6, but can also work with IPv4. Both versions support AH and ESP. AH provides authentication and integrity using HMAC. ESP provides confidentiality, integrity, and authentication using HMAC and AES or 3DES. IPsec creates secure tunnels for VPNs using UDP port 500 for IKE.

IPv4—Internet Protocol version 4. Identifies hosts using a 32-bit IP address. IPv4 is expressed in dotted decimal format with decimal numbers separated by dots or periods like this: 192.168.1.1.

IPv6—Internet Protocol version 6. Identifies hosts using a 128-bit address. IPv6 has a significantly larger address space than IPv4. IPsec is built in to IPv6 and can encrypt any type of IPv6 traffic.

IR—Incident response. Process of responding to a security incident. Organizations often create an incident response policy that outlines procedures and responsibilities of personnel on incident response teams.

IRC—Internet Relay Chat. A form of real-time Internet text messaging often used with chat sessions. Some botnets have used IRC channels to control zombie computers through a command-and-control server.

IRT—Incident Response Team. A group of experts who respond to security incidents. Also known as CERT, CIRT, or SIRT.

IRP—Incident Response Procedure. Procedures documented in an incident response policy.

ISA—Interconnection Security Agreement. Specifies technical and security requirements for connections between two or more entities. An ISA includes details on planning, establishing, maintaining, and disconnecting a secure connection between two or more entities.

iSCSI—Internet Small Computer System Interface. A lower-cost alternative to traditional SANs. It supports sending traditional SCSI commands over an IP network.

ISP—Internet Service Provider. A company that provides Internet access to customers.

ISSO—Information Systems Security Officer. A job role within an organization focused on information security.

IT—Information technology. Computer systems and networks used within organizations.

ITCP—IT contingency plan. Part of risk management. Plan to ensure that IT resources remain available after a security incident, outage, or disaster.

IV—Initialization vector. An IV provides randomization of encryption keys to help ensure that keys are not reused. WEP was susceptible to IV attacks because it used relatively small IVs. In an IV attack, the attacker uses packet injection, increasing the number of packets to analyze, and discovers the encryption key.

JBOD—Just a Bunch of Disks. Disks installed on a computer but not as a RAID.

KDC—Key Distribution Center. Also known as TGT server. Part of the Kerberos protocol used for network authentication. The KDC issues timestamped tickets that expire.

L2TP—Layer 2 Tunneling Protocol. Tunneling protocol used with VPNs. L2TP is commonly used with IPsec (L2TP/IPsec). L2TP uses UDP port 1701. Compare to PPTP, which uses TCP port 1723.

LAN—Local area network. Group of hosts connected within a network.

LANMAN—Local area network manager. Older authentication protocol used to provide backward compatibility to Windows 9x clients. LANMAN passwords are easily cracked due to how they are stored.

LDAP—Lightweight Directory Access Protocol. Language used to communicate with directories such as Microsoft Active Directory. Identifies objects with query strings using codes such as CN=Users and DC=GetCertifiedGetAhead. LDAP uses TCP port 389. Secure LDAP encrypts transmissions with SSL or TLS over TCP port 636. LDAP injection attacks attempt to access or modify data in directory service databases.

LEAP—Lightweight Extensible Authentication Protocol. A modified version of the Challenge Handshake Authentication Protocol (CHAP) created by Cisco. LEAP does not require a digital certificate and Cisco now recommends using stronger protocols such as EAP-TLS.

LSO—Local shared objects or locally shared objects. A Flash cookie created by Adobe Flash player.

MaaS—Monitoring as a Service or Management as a Service. Allows an organization to outsource the management and monitoring of IT resources.

MAC—Mandatory access control. Access control model that uses sensitivity labels assigned to objects (files and folders) and subjects (users). MAC restricts access based on a need-to-know.

MAC—Media access control. A 48-bit address used to identify network interface cards. It is also called a hardware address or a physical address, and is commonly displayed as six pairs of hexadecimal characters. Port security on a switch or an AP can limit access using MAC filtering.

MAC—Message authentication code. Method used to provide integrity for messages. A MAC uses a secret key to encrypt the hash. HMAC is a commonly used version.

Malware—Malicious software. Includes viruses, Trojans, adware, spyware, rootkits, backdoors, logic bombs, and ransomware.

MAN—Metropolitan area network. A computer network that spans a metropolitan area such as a city or a large campus.

MBR—Master Boot Record. An area on a hard disk in its first sector. When the BIOS boots a system, it looks at the MBR for instructions and information on how to boot the disk and load the operating system. Some malware tries to hide here.

MD5—Message Digest 5. A hashing function used to provide integrity. MD5 creates 128-bit hashes, which are also referred to as MD5 checksums. A hash is simply a number created by applying the algorithm to a file or message at different times. Comparing the hashes verifies integrity.

MITM—Man in the middle. A MITM attack is a form of active interception allowing an attacker to intercept traffic and insert malicious code sent to other clients. Kerberos provides mutual authentication and helps prevent MITM attacks.

MOU—Memorandum of understanding. Defines responsibilities of each party, but it is not as strict as an SLA or an ISA. If the parties will be handling sensitive data, they should include an ISA to ensure strict guidelines are in place to protect the data while in transit.

MPLS—Multi-Protocol Layer Switch. A WAN topology provided by some telecommunications companies. Directs data to nodes using labels rather than IP addresses.

MS-CHAP—Microsoft Challenge Handshake Authentication Protocol. Microsoft implementation of CHAP. MS-CHAPv2 provides mutual authentication.

MTBF—Mean time between failures. Provides a measure of a system's reliability and is usually represented in hours. The MTBF identifies the average (the arithmetic mean) time between failures. Higher MTBF numbers indicate a higher reliability of a product or system.

MTTF—Mean time to failure. The length of time you can expect a device to remain in operation before it fails. It is similar to MTBF, but the primary difference is that the MTBF metric indicates you can repair the device after it fails. The MTTF metric indicates that you will not be able to repair a device after it fails.

MTTR—Mean time to recover. Identifies the average (the arithmetic mean) time it takes to restore a failed system. Organizations that have maintenance contracts often specify the MTTR as a part of the contract.

MTU—Maximum Transmission Unit. The MTU identifies the size of data that can be transferred.

NAC—Network access control. Inspects clients for health and can restrict network access to unhealthy clients to a remediation network. Clients run agents and these agents report status to a NAC server. NAC is used for VPN and internal clients. MAC filtering is a form of NAC.

NAT—Network Address Translation. A service that translates public IP addresses to private IP addresses and private IP addresses to public IP addresses. Compare to PAT and DNAT.

NDA—Non-disclosure agreement. Ensures that third parties understand their responsibilities. It is commonly embedded as a clause in a contract with the third party. Most NDAs prohibit sharing data unless you are the data owner.

NDP—Neighbor Discovery Protocol performs several functions on IPv6. For example, it performs functions similar to ARP, which is used on IPv4. It also performs autoconfiguration of device IPv6 addresses and discovers other devices on the network such as the IPv6 address of the default gateway.

NetBIOS—Network Basic Input/Output System (NetBIOS) is a name resolution service for NetBIOS names on internal networks. NetBIOS also includes session services for both TCP and UDP communication. NetBIOS uses UDP ports 137 and 138, and TCP port 139. It can use TCP port 137, but rarely does.

NFC—Near field communication. A group of standards used on mobile devices that allow them to communicate with other nearby mobile devices. Many credit card readers support payments using NFC technologies with a smartphone.

NIC—Network interface card. Provides connectivity to a network.

NIDS—Network-based intrusion detection system. A NIDS is installed on network devices, such as routers or firewalls and monitors network traffic. It can detect network-based attacks.

NIPS—Network-based intrusion prevention system. An IPS that monitors the network. An IPS can actively monitor data streams, detect malicious content, and stop attacks in progress.

NIST—National Institute of Standards and Technology. NIST is a part of the U.S. Department of Commerce, and it includes an Information Technology Laboratory (ITL). The ITL publishes special publications related to security that are freely available for download at <http://csrc.nist.gov/publications/PubsSPs.html>.

NOP—No operation, sometimes listed as NOOP. NOP instructions are often used in a buffer overflow attack. An attacker often writes a large number of

NOP instructions as a NOP sled into memory, followed by malicious code. Some processors use hexadecimal code x90 for NOP so a string of x90 characters indicates a potential buffer overflow attack.

NOS—Network Operating System. Software that runs on a server and enables the server to manage resources on a network.

NoSQL—Not only Structured Query Language. An alternative to traditional SQL databases. NoSQL databases use unstructured query language queries instead of traditional SQL queries.

NTFS—NT File System. A file system used in Microsoft operating systems that provides security. NTFS uses the DAC model.

NTLM—New Technology LANMAN. Authentication protocol intended to improve LANMAN. The LANMAN protocol stores passwords using a hash of the password by first dividing the password into two 7-character blocks, and then converting all lowercase letters to uppercase. This makes LANMAN easy to crack. NTLM stores passwords in LANMAN format for backward compatibility, unless the passwords are greater than 15 characters. NTLMv1 is older and has known vulnerabilities. NTLMv2 is newer and secure.

NTP—Network Time Protocol. Protocol used to synchronize computer times.

OCSP—Online Certificate Status Protocol. An alternative to using a CRL. It allows entities to query a CA with the serial number of a certificate. The CA answers with good, revoked, or unknown.

OLA—Open License Agreement. A volume licensing agreement allowing an organization to install software on multiple systems.

OS—Operating system. Includes Windows, Linux, and Apple iOS systems. OSs are hardened to make them more secure from their default installation.

OSI—Open Systems Interconnection. The OSI reference model conceptually divides different networking requirements into seven separate layers.

OWASP—Open Web Application Security Project. International standard proposed for vulnerability assessment scanners to follow.

P2P—Peer-to-peer. P2P applications allow users to share files such as music, video, and data over the Internet. Data leakage occurs when users install P2P software and unintentionally share files. Organizations often block P2P software at the firewall.

PaaS—Platform as a Service. A cloud computing technology that provides cloud customers with a preconfigured computing platform they can use as needed. PaaS is a fully managed platform, meaning that the vendor keeps the platform up to date with current patches. Compare to IaaS and SaaS.

PAC—Proxy Auto Configuration. Method used to automatically configure systems to use a proxy server.

PAM—Pluggable Authentication Modules. A library of APIs used for authentication-related services.

PAN—Personal area network. A network of devices close to a single person.

PAP—Password Authentication Protocol. An older authentication protocol where passwords or PINs are sent across the network in cleartext. CHAP is more secure. PAP uses PPP.

PAT—Port Address Translation. A form of NAT that translates public IP addresses to private IP addresses, and private IP addresses back to public IP addresses. PAT uses a single public IP address. Compare to DNAT.

PBKDF2—Password-Based Key Derivation Function 2. A key stretching technique that adds additional bits to a password as a salt. This method helps prevent brute force and rainbow table attacks. Bcrypt is a similar key stretching technique.

PBX—Private Branch Exchange. A telephone switch used with telephone calls.

PCAP—Packet Capture. A file that contains packets captured from a protocol analyzer or sniffer.

PDF—Portable Document Format. Type of file for documents. Attackers have embedded malware in PDFs.

PEAP—Protected Extensible Authentication Protocol. PEAP provides an extra layer of protection for EAP and it is sometimes used with 802.1x. PEAP requires a certificate on the 802.1x server. *See also* EAP-TTLS and EAP-TLS.

PED—Personal Electronic Device. Small devices such as cell phones, radios, CD players, DVD players, video cameras, and MP3 players.

PGP—Pretty Good Privacy. Commonly used to secure email communications between two private individuals but is also used in companies. It provides confidentiality, integrity, authentication, and non-repudiation. It can digitally sign and encrypt email. It uses both asymmetric and symmetric encryption.

PII—Personally Identifiable Information. Information about individuals that can be used to trace a person's identity, such as a full name, birth date, biometric data, and identifying numbers such as a Social Security number (SSN). Organizations have an obligation to protect PII and often identify procedures for handling and retaining PII in data policies such as encrypting it.

PIN—Personal identification number. A number known by a user and entered for authentication. PINs are often combined with smart cards to provide dual-factor authentication.

PIV—Personal Identity Verification card. A specialized type of smart card used by U.S. federal agencies. It includes photo identification and provides confidentiality, integrity, authentication, and non-repudiation for the users. It is similar to a CAC.

PKI—Public Key Infrastructure. Group of technologies used to request, create, manage, store, distribute, and revoke digital certificates. Certificates include public keys along with details on the owner of the certificate, and on the CA that issued the certificate. Certificate owners share their public key by sharing a copy of their certificate. A PKI requires a trust model between CAs and most trust models are hierarchical and centralized with a central root CA.

POP3—Post Office Protocol v3. Used to transfer email from mail servers to clients. POP3 uses TCP port 110.

POTS—Plain old telephone service. Voice-grade telephone service using traditional telephone wires.

PPP—Point-to-Point Protocol. Used to create remote access connections. Used by PAP and CHAP.

PPTP—Point-to-Point Tunneling Protocol. Tunneling protocol used with VPNs. PPTP uses TCP port 1723.

PSK—Preshared key. A secret shared among different systems. Wireless networks support Personal mode, where each device uses the same PSK. In contrast, Enterprise mode uses an 802.1x or RADIUS server for authentication.

PTZ—Pan tilt zoom. Refers to cameras that can pan (move left and right), tilt (move up and down), and zoom to get a closer or a wider view.

RA—Recovery agent. A designated individual who can recover or restore cryptographic keys. In the context of a PKI, a recovery agent can recover private keys to access encrypted data, or in some situations, recover the data without recovering the private key. In some cases, recovery agents can recover the private key from a key escrow.

RADIUS—Remote Authentication Dial-In User Service. Provides central authentication for remote access clients. RADIUS uses symmetric encryption to encrypt the password packets and it uses UDP. In contrast, TACACS+ encrypts the entire authentication process and uses TCP. Diameter is an improvement over RADIUS.

RAID—Redundant array of inexpensive disks. Multiple disks added together to increase performance or provide protection against faults. RAID help prevent disk subsystems from being a single point of failure.

RAID-0—Disk striping. RAID-0 improves performance, but does not provide fault tolerance.

RAID-1—Disk mirroring. RAID-1 uses two disks and provides fault tolerance.

RAID-5—Disk striping with parity. RAID-5 uses three or more disks and provides fault tolerance. It can survive the failure of a single drive.

RAID-6—Disk striping with parity. RAID-6 uses four or more disks and provides fault tolerance. It can survive the failure of two drives.

RAM—Random access memory. Volatile memory within a computer that holds active processes, data, and applications. Data in RAM is lost when the computer is turned off. Memory forensics analyzes data in RAM.

RAS—Remote Access Service. Provides access to an internal network from an outside source location using dial-up or a VPN.

RAT—Remote access tool. Commonly used by APTs and other attackers. A RAT gives an attacker full control over a user's system from a remote location over the Internet.

RC—Ron's Code or Rivest's Cipher. Symmetric encryption algorithm that includes versions RC2, RC4, RC5, and RC6. RC4 is a stream cipher, and RC5 and RC6 are block ciphers.

RC4—Rivest Cipher 4. A popular stream cipher. RC4 was implemented incorrectly in WEP, causing vulnerabilities. A rare spelling for RC4 is RSA Variable Key Size Encryption Algorithm.

RDP—Remote Desktop Protocol. Used to connect to remote systems. Microsoft uses RDP in different services such as Remote Desktop Services and Remote Assistance. RDP uses either port TCP 3389 or UDP 3389.

RFI—Radio frequency interference. Interference from RF sources such as AM or FM transmitters. RFI can be filtered to prevent data interference, and cables can be shielded to protect signals from RFI.

RFID—Radio frequency identification. RFID methods are often used for inventory control.

RIPEMD—RACE Integrity Primitives Evaluation Message Digest. A hash function used for integrity. It creates fixed-length hashes of 128, 160, 256, or 320 bits.

ROI—Return of investment or return on investment. A performance measure used to identify when an investment provides a positive benefit to the investor. It is sometimes considered when evaluating the purchase of new security controls.

Role-BAC—Role-based access control. An access control model that uses roles based on jobs and functions to define access and it is often implemented with groups (providing group-based privileges). Often uses a matrix as a planning document to match roles with the required privileges.

RPO—Recovery point objective. The recovery point objective (RPO) refers to the amount of data you can afford to lose by identifying a point in time where data loss is acceptable. It is related to RTO and the BIA often includes both RTOs and RPOs.

RSA—Rivest, Shamir, and Adleman. An asymmetric algorithm used to encrypt data and digitally sign transmissions. It is named after its creators, Rivest, Shamir, and Adleman. RSA uses both a public key and a private key in a matched pair.

RSTP—Rapid Spanning Tree Protocol. An improvement over STP. STP and RSTP protocols are enabled on most switches and protect against switching loops, such as those caused when two ports of a switch are connected together.

RTO—Recovery time objective. An RTO identifies the maximum amount of time it should take to restore a system after an outage. It is derived from the maximum allowable outage time identified in the BIA.

RTP—Real-time Transport Protocol. A standard used for delivering audio and video over an IP network.

Rule-BAC—Rule-based access control. An access control model that uses rules to define access. Rule-based access control is based on a set of approved instructions, such as an access control list, or rules that trigger in response to an event such as modifying ACLs after detecting an attack.

S/MIME—Secure/Multipurpose Internet Mail Extensions. Used to secure email. S/MIME provides confidentiality, integrity, authentication, and non-repudiation. It can digitally sign and encrypt email, including the encryption of email at rest and in transit. It uses RSA, with public and private keys for encryption and decryption, and depends on a PKI for certificates.

SaaS—Software as a Service. A cloud computing technology that provides applications over the Internet. Web mail is an example of a cloud-based

technology. Compare to IaaS and PaaS.

SAML—Security Assertions Markup Language. An XML-based standard used to exchange authentication and authorization information between different parties. SAML provides SSO for web-based applications.

SAN—Storage Area Network. A specialized network of high-speed storage devices.

SCADA—Supervisory control and data acquisition. Typically industrial control systems within large facilities such as power plants or water treatment facilities. SCADA systems are often contained within isolated networks that do not have access to the Internet, but are still protected with redundant and diverse security controls. SCADA systems can be protected with NIPS systems and VLANs.

SCAP—Security Content Automation Protocol. A set of security specifications for various applications and operating systems. Compliance tools such as vulnerability scanners use these to check systems for compliance.

SCEP—Simple Certificate Enrollment Protocol. A method of requesting a certificate from a CA.

SCP—Secure Copy. Based on SSH, SCP allows users to copy encrypted files over a network. SCP uses TCP port 22.

SCSI—Small Computer System Interface. Set of standards used to connect peripherals to computers. Commonly used for SCSI hard disks and/or tape drives.

SDLC—Software Development Life Cycle. A software development process. Many different models are available.

SDLM—Software Development Life Cycle Methodology. The practice of using a SDLC when developing applications.

SEH—Structured Exception Handler. Module within an application that handles errors or exceptions. It prevents applications from crashing or responding to events that can be exploited by attackers.

SELinux—Security-Enhanced Linux. An operating system platform that prevents malicious or suspicious code from executing on both Linux and Unix systems. It is one of the few operating systems that use the MAC model.

SFTP—Secure File Transfer Protocol. An extension of Secure Shell (SSH) using SSH to transmit the files in an encrypted format. SFTP transmits data using TCP port 22.

SHA—Secure Hash Algorithm. A hashing function used to provide integrity. SHA-1 uses 160 bits, and SHA-256 uses 256 bits. As with other hashing algorithms, SHA verifies integrity.

SHTTP—Secure Hypertext Transfer Protocol. An alternative to HTTPS. Rarely used.

SID—Security identifier. Unique set of numbers and letters used to identify each user and each group in Microsoft environments.

SIEM—Security Information and Event Management. A security system that attempts to look at security events throughout the organization.

SIM—Subscriber Identity Module. A small smart card that contains programming and information for small devices such as cell phones.

SIRT—Security Incident Response Team. A group of experts who respond to security incidents. Also known as CERT, CIRT, or IRT.

SLA—Service level agreement. An agreement between a company and a vendor that stipulates performance expectations, such as minimum uptime and maximum downtime levels. Organizations use SLAs when contracting services from service providers such as Internet Service Providers (ISPs).

SLE—Single loss expectancy. The SLE identifies the amount of each loss and is used to measure risk with ALE and ARO in a quantitative risk assessment. The calculation is $SLE \times ARO = ALE$.

SMTP—Simple Mail Transfer Protocol. Used to transfer email between clients and servers and between email servers and other email servers. SMTP uses TCP port 25.

SNMP—Simple Network Management Protocol. Used to manage and monitor network devices such as routers or switches. SNMP agents report information via notifications known as SNMP traps, or SNMP device traps. SNMP uses UDP ports 161 and 162.

SONET—Synchronous Optical Network Technologies. A multiplexing protocol used to transfer data over fiber-optic cable.

SPIM—Spam over Internet Messaging. A form of spam using instant messaging that targets instant messaging users.

SPOF—Single point of failure. An SPOF is any component whose failure results in the failure of an entire system. Elements such as RAID, failover clustering, UPS, and generators remove many single points of failure.

SQL—Structured Query Language. Used by SQL-based databases, such as Microsoft SQL Server. Web sites integrated with a SQL database are subject to SQL injection attacks. Input validation with forms and stored procedures help prevent SQL injection attacks. Microsoft SQL Server uses TCP port 1433 by default.

SSD—Solid State Drive. A drive used in place of a traditional hard drive. An SSD has no moving parts, but instead stores the contents as nonvolatile memory. SSDs are much quicker than traditional drives.

SSH—Secure Shell. SSH encrypts a wide variety of traffic such as SCP, SFTP, Telnet, and TCP Wrappers. SSH uses TCP port 22. SSH is a more secure alternative than Telnet.

SSID—Service Set Identifier. Identifies the name of a wireless network. Disabling SSID broadcast can hide the network from casual users, but an attacker can easily discover it with a wireless sniffer. It's recommended to change the SSID from the default name.

SSL—Secure Sockets Layer. Used to encrypt data in transit with the use of certificates. SSL is used with HTTPS to encrypt HTTP traffic and can also encrypt SMTP and LDAP traffic.

SSO—Single sign-on. Authentication method where users can access multiple resources on a network using a single account. SSO can provide central

authentication against a federated database for different operating systems.

SSTP—Secure Socket Tunneling Protocol. A tunneling protocol that encrypts VPN traffic using SSL over TCP port 443.

STP—Shielded twisted-pair. Cable type used in networks that includes shielding to prevent interference from EMI and RFI. It can also prevent data from emanating outside the cable.

STP—Spanning Tree Protocol. Protocol enabled on most switches that protects against switching loops. A switching loop can be caused if two ports of a switch are connected together.**SYN**—Synchronize. The first packet in a TCP handshake. In a SYN flood attack, attackers send this packet, but don't complete the handshake after receiving the SYN/ACK packet. A flood guard is a logical control that protects against SYN flood attacks.

TACACS+—Terminal Access Controller Access-Control System+. Provides central authentication for remote access clients and used as an alternative to RADIUS. TACACS+ uses TCP port 49. It encrypts the entire authentication process, compared with RADIUS, which only encrypts the password. It uses multiple challenges and responses.

TCO—Total cost of ownership. A factor considered when purchasing new products and services. TCO attempts to identify the cost of a product or service over its lifetime.

TCP—Transmission Control Protocol. Provides guaranteed delivery of IP traffic using a three-way handshake.

TCP/IP—Transmission Control Protocol/Internet Protocol. Represents the full suite of protocols used on the Internet and most internal networks.

TFTP—Trivial File Transfer Protocol. Used to transfer small amounts of data with UDP port 69. In contrast, FTP is used to transfer larger files using TCP ports 20 and 21.

TGT—Ticket Granting Ticket. Used with Kerberos. A KDC (or TGT server) issues timestamped tickets that expire after a certain time period.

TKIP—Temporal Key Integrity Protocol. Wireless security protocol introduced to address the problems with WEP. TKIP was used with WPA but has been deprecated. WPA2 with CCMP is recommended instead.

TLS—Transport Layer Security. Used to encrypt data in transit. TLS is the replacement for SSL and like SSL, it uses certificates issued by CAs. PEAP-TLS uses TLS to encrypt the authentication process and PEAP-TLS requires a CA to issue certificates.

TOTP—Time-based One-Time Password. Similar to HOTP, but it uses a timestamp instead of a counter. One-time passwords created with TOTP expire after 30 seconds.

TPM—Trusted Platform Module. A hardware chip on the motherboard included on many newer laptops. A TPM includes a unique RSA asymmetric key, and when first used, creates a storage root key. TPMs generate and store other keys used for encryption, decryption, and authentication. TPM provides full disk encryption.

TSIG—Transaction Signature. A method of securely providing updates to DNS with the use of authentication.

UAT—User Acceptance Testing. One of the last phases of testing an application before its release.

UDP—User Datagram Protocol. Used instead of TCP when guaranteed delivery of each packet is not necessary. UDP uses a best-effort delivery mechanism.

UEFI—Unified Extensible Firmware Interface. A method used to boot some systems and intended to replace Basic Input/Output System (BIOS) firmware.

UPS—Uninterruptible power supply. A battery backup system that provides fault tolerance for power and can protect against power fluctuations. A UPS provides short-term power giving the system enough time to shut down smoothly, or to transfer to generator power. Generators provide long-term power in extended outages.

URI—Uniform Resource Identifier. Used to identify the name of a resource and always includes the protocol such as *http://GetCertifiedGetAhead.com*.

URL—Uniform Resource Locator. A type of URI. Address used to access web resources, such as *http://GetCertifiedGetAhead.com*. Pop-up blockers can include URLs of sites where pop-ups are allowed.

USB—Universal Serial Bus. A serial connection used to connect peripherals such as printers, flash drives, and external hard disk drives. Data on USB drives can be protected against loss of confidentiality with encryption. Attackers have spread malware through Trojans.

UTM—Unified threat management. A security appliance that combined multiple security controls into a single solution. UTM appliances can inspect data streams for malicious content and often include URL filtering, malware inspection, and content inspection components.

UTP—Unshielded twisted-pair. Cable type used in networks that do not have any concerns over EMI, RFI, or cross talk. If these are a concern, STP is used.

VDI—Virtualization Desktop Infrastructure. Virtualization software designed to reproduce a desktop operating system as a virtual machine on a remote server.

VLAN—Virtual local area network. A VLAN separates or segments traffic. A VLAN can logically group several different computers together, or logically separate computers, without regard to their physical location. It is possible to create multiple VLANs with a single switch. You can also create VLANs with virtual switches.

VM—Virtual machine. A virtual system hosted on a physical system. A physical server can host multiple VMs as servers. Virtualization helps reduce the amount of physical equipment required, reducing overall physical security requirements such as HVAC and power.

VoIP—Voice over IP. A group of technologies used to transmit voice over IP networks. Vishing is a form of phishing that sometimes uses VoIP.

VPN—Virtual private network. Provides access to a private network over a public network such as the Internet. VPN concentrators provide VPN access to large groups of users.

VSAN—Virtual Storage Area Network. A lower-cost alternative to traditional SANs.

VTC—Video teleconferencing. A group of interactive telecommunication technologies that allow people in two or more locations to interact with two-way video and audio transmissions.

WAF—Web application firewall. A firewall specifically designed to protect a web application, such as a web server. A WAF inspects the contents of traffic to a web server, can detect malicious content such as code used in a cross-scripting attack, and block it.

WAP—Wireless access point, sometimes called an access point (AP). Provides wireless clients connectivity to a wired network. Most WAPs use an omnidirectional antenna. You can connect two WLANs together using high-gain directional Yagi antennas. Increasing the power level of a WAP increases the wireless coverage of the WAP. Decreasing the power levels decreases the coverage.

WEP—Wired Equivalent Privacy. Original wireless security protocol. Had significant security flaws and was replaced with WPA, and ultimately WPA2. WEP used RC4 incorrectly making it susceptible to IV attacks, especially when the attacker used packet injection techniques.

WIDS—Wireless intrusion detection system. An IDS used for wireless networks.

WIPS—Wireless intrusion prevention system. An IPS used for wireless networks.

WLAN—Wireless local area network. Network connected wirelessly.

WPA—Wi-Fi Protected Access. Replaced WEP as a wireless security protocol without replacing hardware. Originally used TKIP with RC4 and later implementations support AES. Superseded by WPA2. In WPA cracking attacks, attackers capture the four-way authentication handshake and then use a brute force attack to discover the passphrase.

WPA2—Wi-Fi Protected Access II. Security protocol used to protect wireless transmissions. It supports CCMP for encryption, which is based on AES and is stronger than TKIP, which was originally released with WPA. It uses an 802.1x server for authentication in WPA2 Enterprise mode and a preshared key for WPA2 Personal mode, also called WPA2-PSK.

WPS—Wi-Fi Protected Setup. Allowed users to easily configure a wireless network, often by using only a PIN. WPS brute force attacks can discover the PIN.

WTLS—Wireless Transport Layer Security. Used to encrypt traffic for smaller wireless devices.

XML—Extensible Markup Language. Used by many databases for inputting or exporting data. XML uses formatting rules to describe the data.

XSRF—Cross-site request forgery. Attackers use XSRF attacks to trick users into performing actions on web sites, such as making purchases, without their knowledge. In some cases, it allows an attacker to steal cookies and harvest passwords.

XSS—Cross-site scripting. Attackers use XSS to capture user information such as cookies. Input validation techniques on the server-side help prevent XSS attacks by blocking HTML and JavaScript tags. Many sites prevent the use of < and > characters to block cross-site scripting.

XTACACS—Extended Terminal Access Controller Access-Control System. An improvement over TACACS developed by Cisco Systems and proprietary to Cisco systems. TACACS+ is used more commonly.